Instability of backoff protocols with arbitrary arrival rates

Leslie Ann Goldberg and John Lapinskas*

22 October 2022

Abstract

In contention resolution, multiple processors are trying to coordinate to send discrete messages through a shared channel with sharply limited communication. If two processors inadvertently send at the same time, the messages collide and are not transmitted successfully. An important case is acknowledgement-based contention resolution, in which processors cannot listen to the channel at all; all they know is whether or not their own messages have got through. This situation arises frequently in both networking and cloud computing. The most common acknowledgement-based protocols in practice are backoff protocols — variants of binary exponential backoff are used in both Ethernet and TCP/IP, and both Google Drive and AWS instruct their users to implement it to handle busy periods.

In queueing models, where each processor has a queue of messages, stable backoff protocols are already known (Håstad et al., SICOMP 1996). In queue-free models, where each processor has a single message but processors arrive randomly, it is a long-standing conjecture of Aldous (IEEE Trans. Inf. Theory 1987) that no stable backoff protocols exist for any positive arrival rate of processors. Despite exciting recent results for full-sensing protocols which assume far greater listening capabilities of the processors (see e.g. Bender et al. STOC 2020 or Chen et al. PODC 2021), this foundational question remains open; here instability is only known in general when the arrival rate of processors is at least 0.42 (Goldberg et al. SICOMP 2004). We prove Aldous's conjecture for all backoff protocols outside of a tightly-constrained special case using a new domination technique to get around the main difficulty, which is the strong dependencies between messages.

1 Introduction

In the field of contention resolution, multiple processors (sometimes called "stations") are trying to coordinate to send discrete messages (sometimes called "packets") through a shared channel called a multiple access channel. The multiple access channel is not centrally controlled and the processors cannot communicate, except by listening to the channel. The operation of the channel is straightforward. In each (discrete) time step one or more processors might send messages to the channel. If exactly one message is sent then it is delivered successfully and the sender is notified of the success. If multiple messages are sent then they collide and are not transmitted successfully (so they will have to be re-sent later).

We typically view the entire process as a discrete-time Markov chain. At each time step, new messages arrive at processors with rates governed by probability distributions with total rate $\lambda > 0$. After the arrivals, each processor independently chooses whether to send a message through the channel. A *contention*-resolution protocol is a randomised algorithm that the processors use to decide when to send messages to the channel (and when to wait because the channel is too busy!). Our objective is to find a *stable* protocol [14], which is a protocol with the property that the corresponding Markov chain is positive recurrent, implying

^{*}A short version, without the proofs, will appear in the Proceedings of SODA 2023. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission. All data is provided in full in the results section of this paper.

that there is a stationary distribution bounding the likely extent to which messages build up over time. Other objectives include bounding the expected waiting time of messages and maximising the throughput, which is the rate at which messages succeed.

Issues of contention resolution naturally arise when designing networking protocols [1, 2], but it is also relevant to hardware design that enables concurrency [17, 23] and to interaction with cloud computing services [9, 10]. The example of cloud computing will be an instructive one, so we expand on it. Suppose that an unknown number of users is submitting requests to a server which is struggling under the load, with more users arriving over time. The users do not have access to load information from the server and they do not have knowledge of each other — all they know is whether or not their own requests to the server are getting through. The central question of contention resolution is then: how often should users re-send their requests in order to get everyone's requests through as quickly as possible?

There are two main categories of contention resolution protocol — these differ according to the extent to which processors listen to the channel. In *full-sensing protocols*, processors constantly listen to the shared channel obtaining partial information. For example, in addition to learning whether its own sends are successful, a processor may learn on which steps the channel is quiet (with no sends) [26] or it may learn on which steps there are successful sends [6] or it may learn both [21]. While full-sensing protocols are suitable in many settings, there are important settings where they cannot reasonably be implemented (such as the cloud computing example above). In *acknowledgement-based protocols*, the only information that processors receive about the shared channel is whether their own messages get through.

We distinguish between two ways of modelling message arrival. The earliest work in the field focused on queueing models, in which the number N of processors is fixed and each processor maintains a queue of messages to be sent. These models are appropriate for a static network. A particularly simple example is the slotted ALOHA protocol [24], one of the first networking protocols. In this protocol, if there are Nprocessors with non-empty queues, then these processors send independently with probability 1/N. For large N, this is stable if $\lambda < 1/e$. However, it requires the processors to know the value of N. In order to get around this difficulty, Metcalfe and Boggs proposed binary exponential backoff, in which a processor which has already had i unsuccessful attempts to send a given message waits a random amount of time (a geometric random variable with mean 2^i) before again attempting to send. Binary exponential backoff (with some modifications) forms the basis for Ethernet [19] and TCP/IP [5]. For any N, binary exponential backoff is known to be stable for sufficiently small λ [14, 3]. Unfortunately, this value of λ depends on N and binary exponential backoff is unstable if λ is sufficiently large [16]. Remarkably, Håstad, Leighton and Rogoff [16] showed that *polynomial backoff* (where the waiting time after the *i*'th collision is a geometric random variable with expectation i^{α} for some $\alpha > 1$) is stable for all $\lambda \in (0, 1)$. For contention resolution with queues even more powerful full-sensing protocols are known — in particular, there is a stable fullsensing protocol even for the more general model in which some specified pairs of processors are allowed to use the channel simultaneously [25, 26].

In this paper, we focus on *queue-free models*, which allow for dynamic networks and are more appropriate for public wi-fi hotspots [1] or cloud computing [9, 10]. We again consider these models in discrete time. In these models, processors arrive in the system according to a Poisson distribution with rate λ , and each processor only wants to send a single message rather than maintaining a queue; in fact, we typically identify the processors with the messages that they are trying to send. As usual, only one message can pass through the channel at any given time step. In this setting, an acknowledgement-based protocol can be viewed as a joint distribution $(T_1, T_2, ...)$ of times. For each message, the corresponding processor independently samples $(\tau_1, \tau_2, ...)$ from $(T_1, T_2, ...)$. If the message does not get through during the first j - 1 times that it is sent then the processor waits for τ_j time steps before sending it for the j'th time. An important special case is that of *backoff protocols*, in which $(T_1, T_2, ...)$ is a tuple of independent geometric variables. Equivalently, a backoff protocol is associated with a *send sequence* of probabilities $\mathbf{p} = p_0, p_1, ...$ such that, if a processor has already had j unsuccessful sends, then it will send its message on the following time step with probability p_j ; thus $1/p_j$ is the expected waiting time $\mathbb{E}(T_j)$. For example, the case $p_j = 2^{-j}$ gives rise to the binary exponential backoff protocol that we have already described. This protocol is widely used in the queue-free model: both AWS and Google Drive advise users to implement binary exponential backoff when using their services [9, 10].

In the queue-free setting, there has been a great deal of interesting work developing full-sensing protocols and proving that these perform well. Some of this is described in the survey of Chlebus [8]. See also [21, 11, 15]. More recently, Bender et al. [6] have shown that in the full-sensing model without collision detection (where processors listen to the channel to learn on which steps there are successful sends but are unable to distinguish between silence and collisions) there is a full-sensing protocol which achieves constant throughput even when the message arrival is adversarial rather than random. Chen et al. [7] demonstrate a full-sensing protocol that can achieve a decent throughput, even in the presence of jamming. Despite these advances regarding full-sensing protocols, and other protocols assuming more capabilities from processors than acknowledgement-based protocols [22, 13], for acknowledgement-based protocols the most fundamental possible question remains open: **do stable protocols exist at all?** Indeed, this problem remains open even for backoff protocols, and most work on the question has focused on this case.

The following foundational conjecture was made by Aldous [4] in 1987, and is widely believed. It is the focus of this work.

Conjecture 1 (Aldous's Conjecture). *In the queue-free setting, no backoff protocol is stable for any positive value of* λ *.*

Aldous's conjecture remains open to this day. It has been proved for arrival rates $\lambda > 0.42$ [12], but for arbitrary arrival rates the only known results concern special cases which avoid a central difficulty inherent to the problem. Consider a backoff protocol with send sequence p_0, p_1, \ldots and arrival rate $\lambda > 0$. For all integers $j, t \ge 1$, write $b_j(t)$ for the set of messages in the system at time t which have already sent j times (all unsuccessfully). Write $S_i(t)$ for the number of messages in $b_i(t)$ which send at time t, $S_0(t)$ for the number of "newborn" messages which send for the first time at time t, and $S(t) = S_0(t) + S_1(t) + \dots$ for the total number of sends from all messages in the system. Thus a message escapes the system at time t if and only if S(t) = 1, Aldous's conjecture implies that S(t) = 1 for less than a λ proportion of times, i.e. that messages arrive faster than they escape. Very often, the reason that this occurs is that $S(t) \ge 2$ for most values of t. However, it is not hard to show that for all $j \ge 1$ and most times t we have $\mathbb{E}(S_j(t)) = p_j \mathbb{E}(|b_j(t)|) \lesssim \lambda$, and so on most time steps $S_j(t) = 0$; thus to show $S(t) \ge 2$ on most time steps, we must engage with the complicated joint distribution $(|b_1(t)|, |b_2(t)|, ...)$. This is the key difficulty that all current arguments have avoided, which restricts the classes of send sequences to which they apply. The tool that enabled us to prove Aldous's conjecture for most protocols is a new domination technique for bounding this joint distribution. Before stating our result and the new technique we first summarise progress that can be made without engaging with the joint distribution. In the following summary, we classify protocols in terms of the key quantity $1/p_i$, which is the expected waiting time before a message sends after having its *j*'th collision.

- Kelly and MacPhee [18] categorised the class of backoff protocols for which S(t) ≥ 2 for all sufficiently large t. This result covers all protocols with subexponential expected waiting times, i.e. whose send sequences satisfy 1/p_j = o(c^j) as j → ∞ for all c > 1 (see Corollary 9). Since these protocols are unstable in such a very strong way, Kelly and MacPhee are able to avoid working with the joint distribution in favour of applying the Borel-Cantelli lemmas.
- Aldous [4] proved that binary exponential backoff is unstable for all λ > 0, and his argument easily extends to all backoff protocols with exponential expected waiting times, i.e. whose send sequences satisfy 1/p_j = Θ(c^j) as j → ∞ for some c > 1 (see Section 1.2). This proof relies on proving

concentration for specific variables $|b_j(t)|$ as $t \to \infty$, and then applying union bounds over suitable ranges of j and t, again avoiding the joint distribution. This concentration fails in general; for example, if $p_j \ge 3\lambda$, then $b_j(t) = \emptyset$ for most values of t. (See Definition 10 for a more detailed discussion.)

A simple argument known to the authors of [12] (but not previously published) shows that there is no stable backoff protocol which has infinitely many super-exponential expected waiting times, i.e. with an infinite subsequence p_{j1}, p_{j2},... satisfying 1/p_{jk} = ω(c^{jk}) as k → ∞ for all c > 1. We state and prove this as Lemma 7. In this case there is no need to engage with the joint distribution of the S_j(t) variables because the proof relies on bounding S(t) ≥ S₀(t), i.e. only considering newborn messages.

Unfortunately, the above results cannot be combined in any simple way to prove Aldous's conjecture. For example, by including some p_j values such that the expected waiting time $1/p_j$ is less than exponential, it is easy to construct protocols which neither exhibit concentration for specific variables $|b_j(t)|$ nor satisfy $S(t) \ge 2$ for all sufficiently large t (see Section 1.2), so to show that these protocols are unstable we must engage with the joint distribution.

Our main technical contribution is a proof (see Lemma 16) that, roughly speaking, we can dominate the joint distribution of $(|b_1(t)|, |b_2(t)|, ...)$ below by a much simpler collection of independent Poisson variables whenever $\mathbb{E}(S(t)) \to \infty$ as $t \to \infty$. Using this, we are able to almost entirely solve the problem of inconsistent decay rates and prove Aldous's conjecture except in some extreme cases characterised in Definition 4. Before describing these extreme cases, we give some easier-to-state consequences of our main result (Theorem 5). In the following theorems (and throughout the paper) a *backoff process* is a backoff protocol in the queue-free model. The first consequence is that all protocols with monotonically non-increasing send sequences are unstable.

Theorem 2. For every $\lambda \in (0, 1)$ and every monotonically non-increasing send sequence $\mathbf{p} = p_0, p_1, \ldots$, the backoff process with arrival rate λ and send sequence \mathbf{p} is unstable.

We have included Theorem 2 because it has a clean statement, but our proof technique doesn't rely on any kind of monotonicity. For example, our main Theorem, Theorem 5 also has the following corollary.

Theorem 3. Let \mathbf{p} be a send sequence. Let $m_{\mathbf{p}}(n)$ be the median of p_0, \ldots, p_n . Suppose that $m_{\mathbf{p}}(n) = o(1)$. Then for every $\lambda \in (0, 1)$ the backoff process with arrival rate λ and send sequence \mathbf{p} is unstable.

Of course, there is nothing very special about the median. The same would be true of sequences for which any centile is o(1). At this point, we are ready to describe the extreme cases that elude our new proof technique, and to state our main result, which shows instability except in these extreme cases. The extreme cases have the property that the send sequence is almost entirely constant, with occasional exponential (but not super-exponential) waiting times thrown in.

Definition 4. A send sequence \mathbf{p} is LCED ("largely constant with exponential decay") if it satisfies the following properties:

- (*i*) "Largely constant": For all $\eta > 0$, there exists c > 0 such that for infinitely many n, $|\{j \le n : p_j > c\}| \ge (1 \eta)n$.
- (*ii*) "with exponential decay": **p** has an infinite subsequence $(p_{\ell_1}, p_{\ell_2}, ...)$ which satisfies $\log(1/p_{\ell_x}) = \Theta(\ell_x)$ as $x \to \infty$.
- (*iii*) "(but without super-exponential decay)": $\log(1/p_j) = O(j)$ as $j \to \infty$.

As an illustrative example of item (i) taking $\eta = 999/1000$, it implies that as you progress along the send sequence p_1, p_2, \ldots , infinitely often, you will notice that all but $1 - \eta = 0.1\%$ of the p_j 's you have seen so far are bounded below by some constant c. The same holds for values of η that are closer to 1,

but c will be correspondingly smaller. Obviously, this is also true for values of η that are closer to 1 but the corresponding constant c would be smaller. Item (ii) means that there is an infinite subsequence of j's where the expected waiting times (after j failures) is exponentially long. Item (iii) just means that expected waiting times are not more than exponentially long.

With this definition, we can state our main theorem, which proves Aldous's conjecture for all sequences except LCED sequences and extends all previously-known results.

Theorem 5. Let \mathbf{p} be a send sequence which is not LCED. Then for every $\lambda \in (0,1)$ the backoff process with arrival rate λ and send sequence \mathbf{p} is unstable.

As we discuss in Section 1.4, LCED sequences can exhibit qualitatively different behaviour from non-LCED sequences, with arbitrarily long "quiet patches" during which almost every message that sends is successful. Lemma 16, our domination of the joint distribution $(|b_1(t)|, |b_2(t)|, ...)$ by independent Poisson variables, is actually false during these "quiet patches", and so proving Aldous's conjecture for LCED sequences will require new ideas. On a conceptual level, the "quiet patches" exhibited by some LCED sequences are essentially the only remaining obstacle to a full proof of Aldous's conjecture.

Nevertheless, as we show in this paper, our new domination is sufficient to cover all send sequences except the LCED sequences. Thus Lemma 16 makes substantial progress on the long-standing conjecture, the first progress in many years, resulting in strong new instability results such as Theorems 2 and 3.

The remainder of the introduction is structured as follows. Section 1.1 gives the formal definition of a backoff process. Sections 1.2 and 1.3 sketch the proof of Theorem 5, with Section 1.2 giving an overview of the relevant existing proof techniques and Section 1.3 explaining our novel ideas. Section 1.4 discusses the remaining obstacles to proving Conjecture 1.

1.1 Formal definitions

We say that a stochastic process is *stable* if it is positive recurrent, and *unstable* otherwise (i.e. if it is null recurrent or transient). A *backoff process* is a backoff protocol in the queue-free model.

Informally, a backoff process is a discrete-time Markov chain associated with an arrival rate $\lambda \in (0, 1)$ and a send sequence $\mathbf{p} = p_0, p_1, p_2, \ldots$ of real numbers in the range (0, 1]. Following Aldous [4], we identify processors and messages, and we think of these as balls moving through a sequence of bins. Each time a message sends, if no other message sends at the same time step, it leaves the system; otherwise, it moves to the next bin. Thus at time t, the j'th bin contains all messages which have sent j times without getting through (these sends occurred at time steps up to and including time t). The system then evolves as follows at a time step t. First, new messages are added to bin 0 according to a Poisson distribution with rate λ . Second, for all $j \ge 0$, each message in bin j sends independently with probability p_j . Third, if exactly one message sends then it leaves the system, and otherwise all messages that sent from any bin j move to the next bin, bin j + 1.

Remark 6. There is no need to consider arrival rates $\lambda \ge 1$ because it is already known that backoff processes with arrival rate $\lambda \ge 1$ are unstable [12]. We also don't allow $p_j = 0$ since that would trivially cause transience (hence, instability).

Formal definition of backoff processes. A backoff process with arrival rate $\lambda \in (0, 1)$ and send sequence $\mathbf{p} = p_0, p_1, p_2, \dots \in (0, 1]$ is a stochastic process X defined as follows. Time steps t are positive integers. Bins j are non-negative integers. There is an infinite set of balls. We now define the set $b_j^X(t)$, which will be the set of balls in bin j just after (all parts of) the t'th step. Initially, all bins are empty, so for all non-negative integers j, $b_j^X(0) = \emptyset$. For any positive integer t, the t'th step of X involves (i) step initialisation (including birth), (ii) sending, and (iii) adjusting the bins. Step t proceeds as follows.

- Part (i) of step t (step initialisation, including birth): An integer n_t is chosen independently from a Poisson distribution with mean λ. This is the number of *newborns* at time t. The set b₀^X(t) contains the balls in b₀^X(t − 1) together with n_t new balls which are *born* at time t. For each j ≥ 1 we define b_i^X(t) = b_i^X(t − 1).
- Part (ii) of step t (sending): For all j ≥ 0, all balls in b'_j^X(t) send independently with probability p_j. We use send^X(t) for the set of balls that send at time t.
- Part (iii) of step t (adjusting the bins):
 - If $|\mathbf{send}^X(t)| \leq 1$ then any ball in $\mathbf{send}^X(t)$ escapes so for all $j \geq 0$ we define $b_j^X(t) = b_j^{X}(t) \setminus \mathbf{send}^X(t)$.
 - Otherwise, no balls escape but balls that send move to the next bin, so we define $b_0^X(t) = b_0'^X(t) \setminus \operatorname{send}^X(t)$ and, for all $j \ge 1$, $b_j^X(t) = (b_{j-1}'^X(t) \cap \operatorname{send}^X(t)) \cup (b_j'^X(t) \setminus \operatorname{send}^X(t))$.

Finally, we define **balls**^X $(t) = \bigcup_{i} b_{i}^{X}(t)$.

1.2 Technical context

We first formally state the result alluded to in Section 1 which proves instability for backoff protocols whose send sequences decay super-exponentially.

Lemma 7. Let X be a backoff process with arrival rate $\lambda \in (0,1)$ and send sequence $\mathbf{p} = p_0, p_1, \dots$ If, for infinitely many $j, p_j \leq (\lambda p_0/2)^j$, then X is unstable.

We defer the proof to Section 2, but it is simple. Essentially, we dominate the expected time for a newborn ball to leave the process below under the assumption that a sending ball always leaves the process unless another ball is born at the same time. Under the assumptions of Lemma 7, this is infinite.

We next describe the results of Kelly and MacPhee [18] and Aldous [4] in more detail than the previous section. This will allow us in Section 1.3 to clearly identify the regimes in which instability is not known, and to clearly highlight the novel parts of our arguments.

We first introduce a key notion from Aldous [4]. Informally, an *externally-jammed process* is a backoff process in which balls never leave; thus if a single ball sends at a given time step, it moves to the next bin as normal. (See Section 3.1 for a formal definition.) Unlike backoff processes, an externally-jammed process starts in its stationary distribution; thus for all $j \ge 0$, the size of $b_j(0)$ is drawn from a Poisson distribution with mean λ/p_j . There is a natural coupling between a backoff process X and an externally-jammed process Y such that $|b_j^X(t)| \le |b_j^Y(t)|$ for all j and t (see Observation 22); thus an externally-jammed process can be used to dominate (from above) the number of balls in a backoff process.

As discussed earlier, Kelly and MacPhee [18] gives a necessary and sufficient condition for infinitely many messages to get through. In our context, the relevant case of their result can be stated as follows. Given a send sequence **p**, let W_0, W_1, \ldots be independent geometric variables such that W_j has parameter p_j for all j. Then for all $\tau \ge 0$, we define

$$\mu_{\tau}(\mathbf{p}) = \sum_{j=0}^{\infty} \mathbb{P}\Big(\sum_{k=0}^{j} W_k \le \tau\Big).$$

Thus if a ball is born at time t in an externally-jammed process, $\mu_{\tau}(\mathbf{p})$ is the expected number of times that ball sends up to time $t + \tau$.

Theorem 8 ([18, Theorem 3.10]). Let **p** be a send sequence, and suppose that for all $\lambda \in (0, 1)$,

$$\sum_{\tau=0}^{\infty} \mu_{\tau}(\mathbf{p}) e^{-\lambda \mu_{\tau}(\mathbf{p})} < \infty.$$

Then for all $\lambda \in (0, 1)$, the backoff process X with arrival rate λ and send sequence **p** is unstable. Moreover, with probability 1, only finitely many balls leave X.

The following corollary is proved in Section 2.

Corollary 9. Let **p** be a send sequence such that $\log(1/p_j) = o(j)$ as $j \to \infty$. Then for all $\lambda \in (0, 1)$, the backoff process X with arrival rate λ and send sequence **p** is unstable.

The result that is proved in Aldous [4] is the instability of binary exponential backoff, in which $p_j = 2^{-j}$ for all *j*, for all arrival rates $\lambda > 0$. Aldous's paper says "without checking the details, [he] believe[s] the argument could be modified to show instability for [all backoff protocols]". Unfortunately, this turns out not to be accurate. However, the proof does generalise in a natural way to cover a broader class of backoff protocols. Our own result extends this much further building on new ideas, but it is nevertheless based on a similar underlying framework, so we now discuss Aldous's proof in more detail.

A key notion is that of *noise*. Informally, given a state $\mathbf{x}(t)$ of a backoff process at time t, the noise (Definition 24) of $\mathbf{x}(t)$ is given by

$$f(\mathbf{x}(t)) = \lambda p_0 + \sum_{j=0}^{\infty} p_j |b_j^X(t)|.$$

Thus, the noise of a backoff process at time t is the expected number of sends at time t + 1 conditioned on the state at time t. Unsurprisingly, while the noise is large, multiple balls are very likely to send at each time step so balls leave the process very slowly (see Lemma 44).

A slightly generalised version of Aldous's proof works as follows. The first step is to choose j_0 to be a suitably large integer, and wait until a time t_0 at which bins $1, \ldots, j_0$ are all "full" in the sense that $|b_j(t_0)| \ge c\lambda/p_j$ for some constant c and all $j \le j_0$. (Such a time t_0 exists with probability 1. Observe that in the stationary distribution of an externally-jammed process, bin j contains λ/p_j balls in expectation.) We then define τ_1, τ_2, \ldots with $\tau_\ell = C \sum_{j=0}^{j_0+\ell} (1/p_j)$ for some constant C; thus τ_ℓ is C times the expected number of time steps required for a newborn ball to reach bin $b_{j_0+\ell}$ in a jammed channel. The key step of the proof is then to argue that with suitably low failure probability, for all ℓ and all t satisfying $t_0 + \tau_\ell \le t \le t_0 + \tau_{\ell+1}$, all bins j with $(j_0 + \ell)/10 \le j \le j_0 + \ell$ satisfy $|b_j(t)| \ge \zeta \lambda/p_j$ for some constant ζ . In other words, we prove that with suitably low failure probability, there is a slowly-advancing frontier of bins which are always full; in particular the noise increases to infinity, and the process is transient and hence unstable.

In order to accomplish this key step of the proof, the argument is split into two cases. If t is close to t_0 , there is a simple argument based on the idea that bins 1 through j_0 were full at time t_0 and, since j_0 is large, they have not yet had time to fully empty. For the second case, suppose that t is significantly larger than t_0 , and the goal is (for example) to show that bin j is very likely to be full at time $t = t_0 + \tau_{\ell}$. By structuring the events carefully, one may assume that the process still has large noise up to time t - 1, so with high probability not many balls leave the process in time steps $\{t_0 + 1, \ldots, t\}$. At this point, Aldous uses a time-reversal argument to show that, in the externally-jammed process, bin j fills with balls that are born after time t_0 during an interval of τ_{ℓ} time steps. Under a natural coupling, these balls follow the same trajectory in both the backoff process and the externally-jammed process unless they leave the backoff process; thus, by a union bound, with high probability most of these balls are present in bin j at time t as required. Aldous then applies a union bound over all bins j with $(j_0 + \ell)/10 \le j \le j_0 + \ell$; crucially, the probability bounds

for each individual bin j are strong enough that he does not need to engage with the more complicated joint distribution of the contents of the bins.

Since Aldous works only with binary exponential backoff in [4], he takes $\tau_{\ell} \sim 2^{j_0+\ell}$; his conjecture that his proof can be generalised to all backoff protocols is based on the idea that the definition of τ_{ℓ} could be modified, which gives rise to the more general argument above. Unfortunately, there are very broad classes of backoff protocols to which this generalisation cannot apply. We now define the collection of send sequences which this modified version of Aldous's proof could plausibly handle.

Definition 10. A send sequence \mathbf{p} is reliable if it has the following property. Let $\lambda > 0$, and let X be a backoff process with arrival rate λ and send sequence \mathbf{p} . Then with positive probability, there exists $\zeta > 0$, and times t_0, t_1, t_2, \ldots and collections $\mathcal{B}_0, \mathcal{B}_1, \mathcal{B}_2, \ldots$ of bins such that:

- for all *i*, for all *t* satisfying $t_i \leq t \leq t_{i+1}$, for all $j \in \mathcal{B}_i$, we have $|b_j(t)| \geq \zeta \lambda/p_i$.
- $|\mathcal{B}_i|$ is an increasing sequence with $|\mathcal{B}_i| \to \infty$ as $i \to \infty$.

Indeed, if Aldous's proof works for a send sequence **p**, then it demonstrates that **p** is reliable: Aldous takes $\zeta = 1/2$, $t_i = \tau_i$ for all *i*, and $\mathcal{B}_i = \{(j_0 + i)/10, \dots, j_0 + i\}$ for all *i*. In short, a reliable protocol is one in which, after some "startup time" t_0 , at all times one can point to a large collection of bins which will reliably *all* be full enough to provide significant noise. As discussed above, it is important that they are all full — otherwise it is necessary to delve into the complicated inter-dependencies of the bins.

1.3 Proof sketch

In order to describe the proof of Theorem 5 and explain how it works on unreliable send sequences that are not covered by Aldous's proof techniques, we first set out a guiding example by defining the following family of send sequences.

Example 11. Given $\rho \in (0, 1)$, an increasing sequence **a** of non-negative integers with $a_0 = 0$, and a function $g: \mathbb{N} \to \mathbb{N}$, the (ρ, \mathbf{a}, g) -interleaved send sequence $\mathbf{p} = p_0, p_1, \ldots$ is given by

$$p_{j} = \begin{cases} \rho^{j} & \text{if } a_{2k} \leq j \leq a_{2k+1} - 1 \text{ for some } k \geq 0, \\ g(j) & \text{if } a_{2k+1} \leq j \leq a_{2k+2} - 1 \text{ for some } k \geq 0. \end{cases}$$

Thus the (ρ, \mathbf{a}, g) -interleaved send sequence is an exponentially-decaying send sequence with base ρ spliced together with a second send sequence specified by g, with the splices occurring at points given by \mathbf{a} .

In this section, we will take $g(j) = 1/\log \log j$ and $a_k = 2^{2^k}$. We will refer to the (ρ, \mathbf{a}, g) -interleaved send sequence with this choice of g and \mathbf{a} as a ρ -interleaved send sequence.

Observe that a (ρ, \mathbf{a}, g) -interleaved send sequence fails to be LCED whenever g(j) = o(1). Thus, a ρ -interleaved send sequence is not LCED.

We claim that as long as ρ is sufficiently small, a ρ -interleaved send sequence is not reliable and neither Theorem 8 nor Lemma 7 prove instability; we expand on this claim in Section 1.4. Our main technical result, which we introduce next, will be strong enough to show that backoff protocols with these send sequences are unstable. In order to introduce the technical result we first give some definitions.

Definition 12. *Given an* $\eta \in (0, 1)$ *, let* $\kappa(\eta) = \lceil 3/\eta \rceil$ *.*

Definition 13. Let

$$p_*(\lambda,\eta,\nu) = \min\left\{\frac{\lambda}{200}, \frac{\lambda\eta}{1800\kappa(\eta)^2\log(1/\nu)}\right\}.$$

Definition 14. Fix λ , η and ν in (0, 1). A send sequence $\mathbf{p} = p_0, p_1, \dots$ is (λ, η, ν) -suitable if $p_0 = 1$ and there is a positive integer n_0 such that for all $n \ge n_0$,

- $|\{j \in [n] \mid p_j \leq p_*(\lambda, \eta, \nu)\}| > \eta n$, and
- $\nu^n < p_n$.

Theorem 15. Fix λ , η and ν in (0,1). Let **p** be a (λ, η, ν) -suitable send sequence. Let X be a backoff process with arrival rate λ and send sequence **p**. Then X is transient, and hence unstable.

Theorem 5 will follow from Theorem 15 (actually, from a slightly more technical version of it, stated as Corollary 26), from Corollary 9, and from Lemma 7. We give the proof of Theorem 5, along with the proofs of Theorems 2 and 3, in Section 6. Observe that for all $\lambda, \eta > 0$, every ρ -interleaved send sequence is a $(\lambda, \eta, 2\rho)$ -suitable send sequence, so Theorem 15 does indeed apply.

Fix $\lambda > 0$. Before sketching the proof of Theorem 15, we set out one more piece of terminology. Recall that there is a natural coupling between a backoff process X and an externally-jammed process Y with arrival rate λ and send sequence **p**. Under this coupling, the evolution of X is a deterministic function of the evolution of Y — balls follow the same trajectories in each process, except that they leave X if they send at a time step in which no other ball sends. As such, we work exclusively with Y, calling a ball *stuck* if it is present in X and *unstuck* if it is not. We write **stuck**_j^Y(t) for the set of stuck balls in bin j at time t, and likewise **unstuck**_j^Y(t) for the set of unstuck balls. (See the formal definition of an externally-jammed process in Section 3.1.)

Intuitively, the main obstacle in applying Aldous's proof sketch to a ρ -interleaved send sequence **p** is the long sequences of bins with $p_j = 1/(\log \log j)$. In the externally-jammed process Y, each individual bin j is likely to empty of balls — and hence also of stuck balls — roughly once every $(\log j)^{\lambda}$ steps. (Indeed, in the stationary distribution, bin j is empty with probability $e^{-\lambda/g(j)} = (\log j)^{-\lambda}$.) As such, the individual bins do not provide a consistent source of noise as reliability would require. We must instead engage with the joint distribution of bins and argue that a long sequence of bins with large p_j , taken together, is likely to continue providing noise for a very long time even if individual bins empty. Unfortunately, these bins are far from independent: for example, if $|\mathbf{stuck}_j^Y(t)| = |\mathbf{stuck}_{j+1}^Y(t)| = 0$, then we should also expect $|\mathbf{stuck}_{j+2}^Y(t)| = 0$. By working only with individual reliably-full bins, Aldous was able to sidestep this issue with a simple union bound, but this is not an option for a ρ -interleaved send sequence **p**.

Observe that since Y starts in its stationary distribution, for all $t \ge 0$ the variables $\{|b_j^Y(t)|: j \ge 0\}$ are mutually independent Poisson variables with $\mathbb{E}(|b_j^Y(t)|) = \lambda/p_j$. Intuitively, we might hope that while the noise of Y is high and very few balls are becoming unstuck, most balls in Y will remain stuck from birth and we will have $|\mathbf{stuck}_j^Y(t)| \approx |b_j^Y(t)|$. If this is true, we can hope to dominate the variables $|\mathbf{stuck}_j^Y(t)|$ below by mutually independent Poisson variables as long as the noise remains large. Our largest technical contribution — and the hardest part of the proof — is making this idea rigorous, which we do in Lemma 16. (Strictly speaking, we use a slightly more technical version stated as Lemma 45.) This lemma is the heart of our proof; we spend Section 3 setting out the definitions needed to state it formally, then give the actual proof in Section 4. With Lemma 45 in hand, we then prove Theorem 25 in Section 5, using a version of Aldous's framework in which a union bound over large bins is replaced by Chernoff bounds over a collection of independent Poisson variables which dominate a non-independent collection of smaller bins from below.

For the rest of this section, we will focus on the statement and proof of Lemma 45. A fundamental difficulty is that we can only hope for such a domination to work while the noise of Y is large, but the noise is just a weighted sum of terms $|\mathbf{stuck}_{j}^{Y}(t)|$, which are precisely the variables we are concerned with — we therefore expect a huge amount of dependency. To resolve this, we define the *two-stream process* in Section 3.2. Essentially, we split Y into a pair $T = (Y^A, Y^B)$ of two externally-jammed processes Y^A and Y^B , each with arrival rate $\lambda/2$, in the natural way. We then say that a ball becomes unstuck in Y^A

if it sends on a time step at which no other ball in Y^B sends, and likewise for Y^B . Thus only balls from Y^B can prevent balls in Y^A from becoming unstuck, and only balls from Y^A can prevent balls in Y^B from becoming unstuck. Each stream of the process acts as a relatively independent source of collisions for the other stream while its noise is high. Naturally, there is a simple coupling back to the originally externally-jammed process under which stuck balls in Y are dominated below by stuck balls in T.

We must also explain what we mean by T having "large noise". We give an informal overview here, and the formal definitions in Section 3.5. We define a "startup event" $\mathcal{E}_{init}(t_0)$ which occurs for some t_0 with probability 1, and which guarantees very large noise at time t_0 (governed by a constant C_{init}). We divide the set of bins into blocks B_1, B_2, \ldots of exponentially-growing length, and define a map $\tau \mapsto bins(\tau)$ from times $t_0 + \tau$ to blocks B_i . We define a constant $\zeta > 0$, and say that T is t_0 -jammed for τ if both Y^A and Y^B have at least $\zeta |bins(\tau - 1)|$ noise at time $t_0 + \tau - 1$. (Observe that if every ball were stuck, then in expectation the balls in $bins(\tau - 1)$ would contribute $\lambda |bins(\tau - 1)|$ total noise, so jammedness says that on average the bins in $bins(\tau - 1)$ are "almost full".) We can now state a somewhat simplified version of Lemma 45.

Lemma 16. Fix λ , η and ν in (0,1) and a (λ, η, ν) -suitable send sequence \mathbf{p} with $p_0 = 1$. Let Y^A and Y^B be independent externally-jammed processes with arrival rate $\lambda/2$ and send sequence \mathbf{p} and consider the two-stream externally-jammed process $T = T(Y^A, Y^B)$. Let t_0 and τ be sufficiently large integers. Then there is a coupling of

- T conditioned on $\mathcal{E}_{init}(t_0)$,
- a sample $\{Z_j^A \mid j \in bins(\tau)\}$ where each Z_j^A is chosen independently from a Poisson distribution with mean $\lambda/(4p_j)$, and
- a sample $\{Z_j^B \mid j \in bins(\tau)\}$ where each Z_j^B is chosen from a Poisson distribution with mean $\lambda/(4p_j)$ and these are independent of each other but not of the $\{Z_j^A\}$ values

in such a way that at least one of the following happens:

- T is not t_0 -jammed for τ , or
- for all $j \in bins(\tau)$, $|\mathbf{stuck}_j^T(A, t_0 + \tau)| \ge Z_j^A$ and $|\mathbf{stuck}_j^T(B, t_0 + \tau)| \ge Z_j^B$.

To prove Lemma 16 (or its full version Lemma 45), we first formalise the idea that Y^A can act as a "relatively independent" source of noise to prevent balls in Y^B becoming unstuck (and vice versa). In Section 4.1, we define a *random unsticking process*; this is an externally-jammed process in which balls become unstuck on sending based on independent Bernoulli variables, rather than the behaviour of other balls. In Lemma 51, we dominate the stuck balls of Y^A below by a random unsticking process R for as long as Y^B is "locally jammed", and likewise for the stuck balls of Y^B .

In order to analyse the random unsticking process, we then make use of a time reversal, defining a *reverse* random unsticking process \tilde{R} in Section 4.2 which we couple to R using a probability-preserving bijection between ball trajectories set out in Lemma 47. We then exploit the fact that balls move independently and the fact that the number of balls following any given trajectory is a Poisson random variable to prove Lemma 16, expressing the set of balls in $\operatorname{stuck}_{j}^{\tilde{R}}$ as a sum of independent Poisson variables (one per possible trajectory) with total mean at least $\lambda/(4p_j)$. Unfortunately, the coupling between R and \tilde{R} does not run over all time steps — for example, since we need to condition on $\mathcal{E}_{\operatorname{init}}(t_0)$, it certainly cannot cover times before t_0 in R. As a result, we cannot analyse all trajectories this way; instead, we define a subset $\operatorname{Fill}_{j}^{\tilde{R}}(t)$ of balls in bin j of R at time t which are born suitably late in R, and analyse the corresponding set $\operatorname{Fill}_{j}^{\tilde{R}}(t)$ of balls in \tilde{R} (see

Section 4.6). We then prove in Lemma 48 that the variables $|\operatorname{Fill}_{j}^{\widetilde{R}}(t)|$ for $j \in \operatorname{bins}(\tau - 1)$ are dominated below by a tuple of independent Poisson variables; this bound then propagates back through the couplings to a lower bound on $|\operatorname{Fill}_{j}^{R}(t)|$ in R and finally to the required bounds on $|\operatorname{stuck}^{Y^{A}}(t)|$ and $|\operatorname{stuck}^{Y^{B}}(t)|$ in T given by Lemma 16.

1.4 Future work

It is natural to wonder where the remaining difficulties are in a proof of Aldous's conjecture. LCED sequences are quite a restricted case, and at first glance one might suspect they could be proved unstable by incremental improvements to Theorem 15, Corollary 9, and Lemma 7. Unfortunately, this is not the case — LCED send sequences can exhibit qualitatively different behaviour than the send sequences covered by these results.

Recall the definition of a (ρ, \mathbf{a}, g) -interleaved send sequence from Example 11. Let \mathbf{p} be such a send sequence, and let X be a backoff process with send sequence \mathbf{p} and arrival rate $\lambda \in (0, 1)$. Unless g is exponentially small, Lemma 7 does not apply to \mathbf{p} for any $\lambda < 2\rho$. Also, Theorem 8 (and hence Corollary 9) does not apply to \mathbf{p} for any $\lambda \in (0, 1)$ whenever ρ is sufficiently small and \mathbf{a} grows suitably quickly. To see this, suppose that ρ is small and that \mathbf{a} grows quickly. Recall the definition of $\mu_{\tau}(\mathbf{p})$ from Theorem 8. It is not hard to show (e.g. via union bounds) that $\mu_{\tau}(\mathbf{p}) \leq \log_{1/\rho}(\tau) + O(1)$ for all τ satisfying $a_{2k} \leq \log_{1/\rho}(\tau) \leq a_{2k+1}/3$ for any k; from this, it follows that

$$\sum_{\tau=0}^{\infty} \mu_{\tau}(\mathbf{p}) e^{-\lambda \mu_{\tau}(\mathbf{p})} \ge \sum_{\tau=0}^{\infty} \mu_{\tau}(\mathbf{p}) e^{-\mu_{\tau}(\mathbf{p})} = \infty \text{ for all } \lambda \in (0,1).$$

Since Theorem 8 is both a necessary and a sufficient condition for a backoff process to send only finitely many times (see [18]), we can conclude that the methods of Kelly and MacPhee do not apply here.

As discussed in Section 1.3, **p** is not LCED whenever g(j) = o(1) as $j \to \infty$. It is also easy to see that **p** is LCED whenever $g(j) = \Theta(1)$ and $a_{2k+1} - a_{2k} = o(a_{2k+2} - a_{2k+1})$ as $k \to \infty$, so in this case our own Theorem 5 does not apply. To show that a simple generalisation of Theorem 15 will not suffice, we introduce the following definition.

Definition 17. A send sequence p has quiet periods if it has the following property for some $\lambda \in (0, 1)$. Let X be a backoff process with arrival rate λ and send sequence p. Then with probability 1, there exist infinitely many time steps on which X has noise less than 1.

Recall that the proof of Theorem 15 (as with Aldous's result [4]) relies on proving that the noise of a backoff process increases to infinity over time, so it is unsuitable for send sequences with quiet periods. Moreover, Corollary 9 is based on Kelly and MacPhee's necessary and sufficient condition for a backoff protocol to have only finitely many successful sends (with probability 1). Backoff protocols with quiet periods have infinitely many successful sends, so cannot be proved unstable by the methods of Corollary 9. Finally, Lemma 7 handles many protocols with quiet periods (e.g. $p_j = 2^{-2^j}$), but these protocols are *always* quiet, and have the property that they can be proved unstable by simple domination arguments (assuming that a ball always succeeds when it sends unless another ball is born at the same time). For many send sequences with quiet periods, these assumptions are not true.

The key problem is that if g(j) = 1/2 (say), ρ is sufficiently small, and a grows sufficiently quickly, then p is very likely to have quiet periods. Indeed, in the externally-jammed process, for large values of k all bins $1, \ldots, a_{2k} - 1$ are likely to simultaneously empty (which happens with probability at least roughly $2^{-a_{2k}}$ in the stationary distribution) long before bin a_{2k} fills with stuck balls and begins contributing significant noise (which will take time at least $(1/\rho)^{a_{2k}}$). Heuristically, we would expect p to alternate between "quiet periods" with very low noise and "noisy periods" in which noise first grows rapidly, then stays high for a very long time. Our current methods are incapable of handling this juxtaposition — in particular, since messages can easily escape from the system in quiet periods, any argument based around our ideas would need to bound the frequency and duration of quiet periods.

However, despite this, we believe we have made significant progress towards a proof of Aldous's conjecture even for LCED sequences. Recall from Section 1.3 that the key part of our proof of Theorem 15 is Lemma 45, which gives a way of dominating the number of stuck balls in different bins below by an *independent* tuple of Poisson variables as long as overall noise is large. Importantly, the proof of Lemma 45 doesn't actually rely on the fact that **p** is suitable. Suitability is important to determining the parameters of the block construction, which is vital in applying Lemma 45 to prove Theorem 15, but the Lemma 45 proof goes through regardless of what these parameters are. As such, we believe a slightly more general version of Lemma 45 will be important even for LCED sequences as a tool for analysing "noisy periods". The main remaining challenge is in showing that such "noisy periods" are common enough to outweigh the effect of "quiet periods" of very low noise.

The remainder of the paper is structured as follows. In Section 2, we set out some standard preliminaries and prove Lemma 7 and Corollary 9. In Section 3 we introduce the main random processes and block construction we will need in order to state Lemma 45, the key lemma for our proof of Theorem 15 (see Section 1.3). In Section 4, we state and prove Lemma 45. In Section 5, we use Lemma 45 to prove Theorem 25. Finally, in Section 6, we use Theorem 25 (together with Lemma 7 and Corollary 9) to prove Theorem 5, and derive Theorems 2 and 3 as corollaries.

2 **Preliminaries**

For all positive integers n, we define [n] to be the set $\{1, 2, ..., n\}$. All logarithms are to the base e unless specified otherwise.

A stochastic process is "stable" if it is positive recurrent. It is "unstable" otherwise. An unstable process can be either null recurrent or transient, We now define the filtration $\mathcal{F}_0^X, \mathcal{F}_1^X, \ldots$ of a backoff process X.

Definition 18. Let X be a backoff process. Since all bins are initially empty, \mathcal{F}_0^X contains no information. For each positive integer t, \mathcal{F}_t^X contains \mathcal{F}_{t-1}^X together with $b_0^{\prime X}(t)$ and, for each $j \ge 0$, $b_j^X(t)$. It is clear that send^X(t) and balls^X(t) can be deduced from \mathcal{F}_t^X .

We now state some common Chernoff bounds.

Lemma 19. (E.g. [20, Theorems 4.5 and 5.4]) Let Ψ be a real random variable with mean μ which is either Poisson or a sum of independent Bernoulli variables. Then for all $0 < \delta < 1$, we have $\mathbb{P}(\Psi \le (1 - \delta)\mu) \le e^{-\delta^2 \mu/2}$.

Lemma 20. Let Ψ be a binomial random variable with mean μ . Then for all x > 1, we have

$$\mathbb{P}(\Psi > x\mu) < e^{-\mu x(\log x - 1)}$$

Proof. By the Chernoff bound of [20, Theorem 4.4(1)], for all $\delta > 0$, we have

$$\mathbb{P}(\Psi \ge (1+\delta)\mu) < \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu}.$$

Taking $x = 1 + \delta$ and putting everything under the exponent implies that for all x > 1,

$$\mathbb{P}(\Psi \ge x\mu) < e^{\mu(x-1) - \mu x \log x} < e^{-\mu x (\log x - 1)},$$

as required.

We will also use some preliminary lemmas. We first prove the following lemma known to to the authors of [12] and stated in Section 1.2.

Lemma 7. Let X be a backoff process with arrival rate $\lambda \in (0,1)$ and send sequence $\mathbf{p} = p_0, p_1, \dots$ If, for infinitely many $j, p_i \leq (\lambda p_0/2)^j$, then X is unstable.

Proof. Suppose that for infinitely many j we have $p_j \leq (\lambda p_0/2)^j$. We will show that X is not positive recurrent by showing that the expected time for X to return to its initial state (the state containing no balls) is infinite.

Let \mathcal{T}_0 be the first time such that $b_1^X(\mathcal{T}_0) \neq \emptyset$, and note that $\mathcal{T}_0 < \infty$ with probability 1. Consider any positive integer t_0 and any value $F_{t_0}^X$ of $\mathcal{F}_{t_0}^X$ such that $\mathcal{T}_0 = t_0$. We will now consider the process Xconditioned on $\mathcal{F}_{t_0}^X = F_{t_0}^X$. Let

 $\mathcal{T} = \min(\{\infty\} \cup \{t > t_0 \mid \mathbf{balls}^X(t) = \emptyset\}).$

Then by the definition of positive recurrence, it suffices to show that $\mathbb{E}[\mathcal{T} \mid \mathcal{F}_{t_0}^X = F_{t_0}^X] = \infty$. Let β be an arbitrary ball in $b_1^X(t_0)$, and let

 $\mathcal{T}' = \min(\{\infty\} \cup \{t > t_0 \mid \beta \in \mathbf{send}^X(t) \text{ and no newborns send at } t\}).$

Observe that β must escape in order for X to return to its initial state, and β cannot escape before time \mathcal{T}' , so $\mathcal{T}' \leq \mathcal{T}$; it therefore suffices to show that $\mathbb{E}[\mathcal{T}' \mid \mathcal{F}_{t_0}^X = F_{t_0}^X] = \infty$. At any given step $t > t_0$, the number of newborns that send is a Poisson random variable with para-

meter λp_0 , so the probability of at least one newborn sending on any given step is $r = 1 - e^{-\lambda p_0}$. We therefore have

$$\mathbb{E}[\mathcal{T}' \mid \mathcal{F}_{t_0}^X = F_{t_0}^X] = \sum_{j \ge 1} r^{j-1} (\frac{1}{p_1} + \dots + \frac{1}{p_j})(1-r) = \frac{1-r}{r} \sum_{j \ge 1} \frac{r^j}{p_j}.$$

By the assumption in the lemma statement, infinitely many values of j contribute a summand which is at least $(2r/(\lambda p_0))^j$ which is at least 1 since $r = 1 - e^{-\lambda p_0} \ge \lambda p_0/2$, thus the sum diverges as required.

Before setting up the structure for our own proof, we also prove Corollary 9, a Corollary of Kelly and MacPhee's Theorem 8.

Corollary 9. Let **p** be a send sequence such that $\log(1/p_j) = o(j)$ as $j \to \infty$. Then for all $\lambda \in (0, 1)$, the backoff process X with arrival rate λ and send sequence **p** is unstable.

Proof. Let p be a send sequence as in the corollary statement. As in Theorem 8, let W_0, W_1, \ldots be independent geometric variables such that W_j has parameter p_j for all j. For all $\tau \ge 2/p_0$, let

$$\mu_{\tau}(\mathbf{p}) = \sum_{j=0}^{\infty} \mathbb{P}\Big(\sum_{k=0}^{j} W_k \le \tau\Big),\tag{1}$$

and let

$$M(\tau) = \max\left\{j \ge 0: \sum_{k=0}^{j} (1/p_k) \le \tau/2\right\}.$$

In order to apply Theorem 8, we first prove that for all $\tau \ge 2/p_0$, $\mu_{\tau}(\mathbf{p}) \ge M(\tau)/2$. By the definition of $M(\tau)$, for all $j \leq M(\tau)$ we have $\mathbb{E}(\sum_{k=0}^{j} W_k) \leq \tau/2$. It follows by Markov's inequality that for all $j \leq M(\tau),$

$$\mathbb{P}\Big(\sum_{k=0}^{j} W_k \le \tau\Big) = 1 - \mathbb{P}\Big(\sum_{k=0}^{j} W_k > \tau\Big) \ge 1/2.$$

It now follows from (1) that

$$\mu_{\tau}(\mathbf{p}) \ge \sum_{j=0}^{M(\tau)} \mathbb{P}\left(\sum_{k=0}^{j} W_k \le \tau\right) \ge M(\tau)/2.$$
(2)

We now prove that $M(\tau) = \omega(\log \tau)$ as $\tau \to \infty$. Let C > 1 be arbitrary; we will show that when τ is sufficiently large, we have

$$\sum_{k=0}^{\lceil C \log \tau \rceil} (1/p_k) \le \tau/2,$$

and hence that $M(\tau) \ge C \log \tau$ as required. Recall that $\log(1/p_j) = o(j)$ as $j \to \infty$, and let j_C be such that $1/p_j \le e^{j/C}/(Ce^4)$ for all $j \ge j_C$. Suppose $\tau \ge 2/p_0$ is large enough that

$$\sum_{k=0}^{j_C-1} (1/p_k) \le \tau/4.$$

We then have

$$\sum_{k=0}^{\lceil C\log\tau\rceil} (1/p_k) \le \frac{\tau}{4} + \sum_{k=j_C}^{\lceil C\log\tau\rceil} (1/p_k) \le \frac{\tau}{4} + \frac{1}{Ce^4} \sum_{k=0}^{\lceil C\log\tau\rceil} e^{k/C} < \frac{\tau}{4} + \frac{1}{Ce^4} \cdot \frac{e^{(1+\lceil C\log\tau\rceil)/C}}{e^{1/C} - 1} + \frac{1}{Ce^4} \cdot \frac{1}{Ce^4}$$

Since $e^x \ge 1 + x$ for all $x \ge 0$, $C(e^{1/C} - 1) \ge 1$ and since C > 1 it follows that

$$\sum_{k=0}^{\lceil C \log \tau \rceil} (1/p_k) \le \frac{\tau}{4} + \frac{1}{e^4} \cdot e^{2 + \log \tau} < \frac{\tau}{2}.$$

Hence $M(\tau) \ge C \log \tau$, so $M(\tau) = \omega(\log \tau)$ as $\tau \to \infty$, as claimed; by (2), it follows that $\mu_{\tau}(\mathbf{p}) = \omega(\log \tau)$ as $\tau \to \infty$.

Finally, we observe that, since $\mu_{\tau}(\mathbf{p}) = \omega(\log \tau)$, for all $\lambda \in (0, 1)$ there is a τ_{λ} such that, for $\tau \geq \tau_{\lambda}$, $\mu_{\tau}(\mathbf{p}) \geq (2/\lambda) \log \tau$. Thus

$$\sum_{\tau=0}^{\infty} \mu_{\tau}(\mathbf{p}) e^{-\lambda \mu_{\tau}(\mathbf{p})} \le \Theta(1) + \sum_{\tau=\tau_{\lambda}}^{\infty} \frac{(2/\lambda) \log \tau}{\tau^2} < \infty.$$

The result therefore follows immediately from Theorem 8.

3 Preparation for our Technical Result

We will prove Theorem 15 by studying stochastic processes that are related to, but not exactly the same as, backoff processes. To do this, it will be useful to assume that the send sequence \mathbf{p} has $p_0 = 1$. The next observation sets up the machinery that enables this.

Observation 21. Fix $\lambda \in (0, 1)$ Let \mathbf{p} be a send sequence. Let \mathbf{p}' be the send sequence that is identical to \mathbf{p} except that $\mathbf{p}'_0 = 1$. Let X be a backoff process with arrival rate λ and send sequence \mathbf{p} . Let X' be a backoff process with arrival rate λp_0 and send sequence \mathbf{p}' . Then there is a coupling of X and X' such that, for every positive integer t and every positive integer j, $b_i^{X'}(t) \supseteq b_i^{X'}(t)$.

Indeed, in defining the coupling, we can identify the newborns of X' with the newborns of X which immediately send.

3.1 Definition of externally-jammed process

In order to study backoff processes it will be helpful to define another process that is associated with an arrival rate $\lambda \in (0,1)$ and a send sequence $\mathbf{p} = p_0, p_1, p_2, \ldots$ of real numbers in the range (0,1] with $p_0 = 1$.

Note that in a backoff process X with $p_0 = 1$ every bin $b_0^X(t)$ is empty, and this is why bins are positive integers j in the following definition.

Y is an externally-jammed process with arrival rate $\lambda \in (0, 1)$ and send sequence **p** if it behaves as follows. Time steps t are positive integers. Bins j are positive integers. We reserve j as an index for bins, and use ℓ when we need a second index. Initialisation: For every positive integer j, an integer x_j is chosen independently from a Poisson distribution with mean λ/p_j . The set $b_j^Y(0)$ is the set containing x_j balls which are "born in process Y at time 0 in bin j". Formally, the names of these balls are $(Y, 0, j, 1), \ldots, (Y, 0, j, x_j)$. For any positive integer t, the t'th step of Y involves (i) step initialisation (including birth), (ii) sending, and (iii) adjusting the bins. The set $b_j^Y(t)$ is the set of balls in bin j just after (all parts of) the t'th step. Step t proceeds as follows.

- Part (i) of step t (step initialisation, including birth): An integer n_t is chosen independently from a Poisson distribution with mean λ . The set $b_0^{\prime Y}(t)$ of *newborns* at time t is the set containing n_t balls which are "born in Y at time t". Formally, the names of these balls are $(Y, t, 1), \ldots, (Y, t, n_t)$. Then for all $j \ge 1, b_i^{\prime Y}(t) = b_i^Y(t-1)$.
- Part (ii) of step t (sending): For all j ≥ 0, all balls in b^Y_j(t) send independently with probability p_j. We use send^Y(t) for the set of balls that send at time t.
- Part (iii) of step t (adjusting the bins): For all $j \ge 1$,

$$b_j^Y(t) = (b_{j-1}'^Y(t) \cap \operatorname{send}^Y(t)) \cup (b_j'^Y(t) \setminus \operatorname{send}^Y(t)).$$

Note that for any distinct j and j', $b_j^Y(t)$ is disjoint from $b_{j'}^Y(t)$. Also, since $p_0 = 1$, $\bigcup_{j \ge 1} b_j^Y(t) = \bigcup_{j \ge 1} b_j^{Y'}(t)$. We now define the filtration $\mathcal{F}_0^Y, \mathcal{F}_1^Y, \ldots$ for the externally-jammed process Y. \mathcal{F}_0^Y contains, for each $j \ge 1$, $b_j^Y(0)$. Then for each positive integer t, \mathcal{F}_t^Y contains \mathcal{F}_{t-1}^Y together with $b_0^{Y'}(t)$ and, for each $j \ge 1$, $b_j^Y(t)$. It is clear that send^Y(t) can be deduced from \mathcal{F}_t^Y . We define balls^Y(t) = $\bigcup_j b_j^Y(t)$. This can also be deduced from \mathcal{F}_t^Y . We will use the following notation: Each ball $\beta = (Y, 0, j, x)$ has birth^Y(β) = t.

It is going to be convenient in our analysis to think about the balls of an externally-jammed process Y as having two states, "stuck" and "unstuck". The set $\mathbf{stuck}^{Y}(t)$ will be the set of all balls in $\mathbf{balls}^{Y}(t)$ that are stuck in the process Y at time t and $\mathbf{unstuck}^{Y}(t) = \mathbf{balls}^{Y}(t) \setminus \mathbf{stuck}^{Y}(t)$. (Given this identity, we will define either $\mathbf{unstuck}^{Y}(t)$ or $\mathbf{stuck}^{Y}(t)$ for each t, the other can be deduced.) We will assign the states "stuck" and "unstuck" in such a way that $\mathbf{stuck}^{Y}(t)$ and $\mathbf{unstuck}^{Y}(t)$ can be deduced from \mathcal{F}_{t}^{Y} , so the assignment of states to balls is just a convenience.

First, $stuck^{Y}(0) = \emptyset$. Then for any $t \ge 1$ we have the following definition.

- $\mathbf{newstuck}^{Y}(t) = \mathbf{stuck}^{Y}(t-1) \cup b_{0}^{Y}(t).$
- stucksend^Y(t) = newstuck^Y(t) \cap send^Y(t).
- If $|\mathbf{stucksend}^{Y}(t)| = 1$ then $\mathbf{unstuck}^{Y}(t) = \mathbf{unstuck}^{Y}(t-1) \cup \mathbf{stucksend}^{Y}(t)$. Otherwise, $\mathbf{unstuck}^{Y}(t) = \mathbf{unstuck}^{Y}(t-1)$.

For $t \ge 1$ we define $\operatorname{stuck}_{j}^{Y}(t) = \operatorname{stuck}^{Y}(t) \cap b_{j}^{Y}(t)$. Note that the stuck balls in an externally-jammed process correspond to a backoff process. This is captured in the following observation.

Observation 22. For any backoff process X with arrival rate $\lambda \in (0, 1)$ and send sequence $\mathbf{p} = p_0, p_1, \ldots$ with $p_0 = 1$, there is a coupling of X with an externally-jammed process Y with arrival rate λ and send sequence **p** such that for all $j \ge 1$ and $t \ge 0$, $|b_i^X(t)| = |\mathbf{stuck}_i^Y(t)|$.

3.2 A two-stream externally-jammed process

Observation 22 shows that in order to study backoff processes we can instead study externally-jammed processes and focus on the stuck balls.

Typically, it is difficult to study externally-jammed processes due to correlation between balls. To alleviate this, we define another related process, called a two-stream externally-jammed process. The idea is that the balls are (arbitrarily) divided into two streams. Keeping track of separate streams enables domination by processes with more independence.

Here is the definition. We build a two-stream externally-jammed process $T = T(Y^A, Y^B)$ with arrival rate λ and send sequence **p** by combining two independent externally-jammed processes Y^A and Y^B , each with arrival rate $\lambda/2$ and send sequence p. The balls of T are the balls of Y^A together with the balls of Y^B . We refer to the former as A-balls of T and to the latter as B-balls of T. For all $t \ge 0$ and $j \ge 1$, $b_j^T(t) = b_j^{Y^A}(t) \cup b_j^{Y^B}(t)$. For all $t \ge 1$ and $j \ge 0$, $b_j'^T(t) = b_j'^{Y^A}(t) \cup b_j'^{Y^B}(t)$. For all $t \ge 1$, $\mathbf{send}^{T}(t) = \mathbf{send}^{Y^{A}}(t) \cup \mathbf{send}^{Y^{B}}(t).$ We define $\mathbf{balls}^{T}(t) = \bigcup_{j} b_{j}^{T}(t)$. For $C \in \{A, B\}$, $\mathbf{balls}^{T}(C, t)$ is the set of all C-balls in $\mathbf{balls}^{T}(t)$.

Also, each C-ball β of T has $\operatorname{birth}^T(\beta) = \operatorname{birth}^{Y^C}(\beta)$.

Lemma 23 below contains the following easy observation: Let Y^A and Y^B be independent externallyjammed processes with arrival rate $\lambda/2$ and consider the two-stream externally-jammed process T = $T(Y^A, Y^B)$. Let Y be an externally-jammed process with arrival rate λ . There is a coupling of Y with T which provides, for every non-negative integer t, a bijection π_t from balls^Y(t) to balls^T(t) such that $\beta \in b_i^Y(t)$ iff $\pi_t(\beta) \in b_i^T(t)$. Given this, it may seem unclear what the point is of defining the two-stream externally jammed process, since this is essentially "the same" as an externally jammed process. The only difference is the names of the balls! The point of it is that we will assign the states stuck and unstuck to balls of T in a different, and more useful, way.

For $C \in \{A, B\}$, we will now define the set $\mathbf{stuck}^T(C, t)$, which will be the set of all balls in **balls**^T(C,t) that are *stuck* in the process T at time t. We also define $stuck^{T}(t) = stuck^{T}(A,t) \cup$ stuck^T(B,t). Similarly to externally jammed processes, it will always be the case that unstuck^T(C,t) =**balls**^T(C,t) \ **stuck**^T(C,t), so it is only necessary to define **stuck**^T(C,t) or **unstuck**^T(C,t) and the other is defined implicitly.

For any $C \in \{A, B\}$ we will define $\mathbf{stuck}^T(C, 0) = \emptyset$. Then for any $t \ge 1$ we have the following definitions:

- For $C \in \{A, B\}$, newstuck^T $(C, t) = \operatorname{stuck}^{T}(C, t-1) \cup b_0^{Y^C}(t)$.
- For $C \in \{A, B\}$, stucksend^T(C, t) = newstuck^T(C, t) \cap send^T(t).
- Case 1: If, for distinct C and C' in $\{A, B\}$, stucksend^T(C, t) = \emptyset and stucksend^T(C', t) $\neq \emptyset$, then
 - β is chosen uniformly at random from stucksend^T(C', t) and
 - unstuck^T(C', t) = unstuck^T(C', t 1) \cup \{\beta\} and
 - $\mathbf{unstuck}^T(C, t) = \mathbf{unstuck}^T(C, t-1).$
- Case 2: Otherwise, for each $C \in \{A, B\}$ unstuck $^{T}(C, t) =$ unstuck $^{T}(C, t-1)$.

Although the sets $b_j^T(t)$, $b_j'^T(t)$, $\operatorname{send}^T(t)$, and $\operatorname{balls}^T(C, t)$ can be deduced from $\mathcal{F}_t^{Y^A}$ and $\mathcal{F}_t^{Y^B}$, the sets $\operatorname{stuck}^T(C, t)$ cannot. This is because $\mathcal{F}_t^{Y^A}$ and $\mathcal{F}_t^{Y^B}$ do not capture any information about the choice of the random ball β (from $\operatorname{stucksend}^T(A, t)$ or $\operatorname{stucksend}^T(B, t)$) that may become unstuck at time t. Thus, we take \mathcal{F}_t^T to be $\mathcal{F}_t^{Y^A}$, $\mathcal{F}_t^{Y^B}$ and the sets $\operatorname{stuck}^T(A, t')$ and $\operatorname{stuck}^T(B, t')$ for $t' \leq t$ so that all information about the evolution of the process T, including the assignment of states to the balls, is captured.

For $C \in \{A, B\}$, $t \ge 1$ and $j \ge 1$, define $\operatorname{stuck}_j^T(C, t) = \operatorname{stuck}^T(C, t) \cap b_j^T(t)$.

We can now couple an externally-jammed process Y to a two-stream externally-jammed process T in the following lemma.

Lemma 23. Let Y^A and Y^B be independent externally-jammed processes with arrival rate $\lambda/2$ and consider the two-stream externally-jammed process $T = T(Y^A, Y^B)$. Let Y be an externally-jammed process with arrival rate λ . There is a coupling of Y with T which provides, for every non-negative integer t, a bijection π_t from **balls**^Y(t) to **balls**^T(t) such that the following two properties hold.

- $\beta \in b_i^Y(t)$ iff $\pi_t(\beta) \in b_i^T(t)$, and
- $\beta \in \mathbf{unstuck}^Y(t)$ implies $\pi_t(\beta) \in \mathbf{unstuck}^T(t)$.

Proof. We first define the bijections π_t . In fact, it is easier to define the inverses π_t^{-1} .

In the initialisation of T, for each $C \in \{A, B\}$, let x_j^C denote the number of C-balls born in T at time 0 in bin j. Then x_j , the number of balls born in Y at time 0 in bin j, is given by $x_j = x_j^A + x_j^B$. The bijection π_0^{-1} from **balls**^T(0) to **balls**^Y(0) then maps the ball $(Y^A, 0, j, x)$ to (Y, 0, j, x) and the ball $(Y^B, 0, j, x)$ to $(Y, 0, j, x_j^A + x)$ (for each x).

For any $t \ge 1$, the bijection π_t^{-1} from **balls**^T(t) to **balls**^Y(t) is the same as π_{t-1} except that it maps the balls in $b'_0^{Y^A}(t) \cup b'_0^{Y^B}(t)$ to $b'_0^Y(t)$ using a similar re-naming of balls.

Note that both properties hold at time t = 0 since all balls are unstuck at time 0.

Suppose, for a positive integer t, that the first t - 1 steps of Y and T have been coupled so that the two properties hold. We will show how to couple the t'th step of the process so that the two properties again hold.

The first property allows us to couple the send decisions so that $\beta \in \text{send}^Y(t)$ iff $\pi_t(\beta) \in \text{send}^T(t)$. This allows us to conclude that the first property holds at time t.

To finish, we wish to show that any ball $\beta \in \mathbf{unstuck}^Y(t)$ has $\pi_t(\beta) \in \mathbf{unstuck}^T(t)$. If $\beta \in \mathbf{unstuck}^Y(t-1)$ then this is immediate by the second property at time t-1; we may therefore focus on the case where β only becomes unstuck at time t, i.e., where $\beta \in \mathbf{stucksend}^Y(t) \cap \mathbf{unstuck}^Y(t)$. By the definition of the Y process, $\beta \in \mathbf{stucksend}^Y(t) \cap \mathbf{unstuck}^Y(t)$ implies that $\mathbf{stucksend}^Y(t) = \{\beta\}$.

By the first property at time t and the second property at time t - 1, if any ball β^* satisfies $\beta^* \in$ stucksend^T(t) then $\pi_t^{-1}(\beta^*) \in$ stucksend^Y(t) so $\pi_t^{-1}($ stucksend^T(t)) \subseteq stucksend^Y(t). Hence |stucksend^T(t) $| \leq 1$. Now we are finished because $\pi_t(\beta)$ is either in unstuck^T(t - 1) (in which case it is also in unstuck^T(t)) or it is in stucksend^T(t) (in which case stucksend^T(t) = { $\pi_t(\beta)$ } so $\pi_t(\beta) \in$ unstuck^T(t)).

3.3 Ball vectors and noise

So far, we have defined backoff processes, externally-jammed processes, and two-stream externally-jammed processes.

Here we collect some notation about these processes that we haven't needed so far, but which will be useful later.

For an externally-jammed process Y we define

$$\mathbf{ballvect}^{Y}(t) = (b_1^{Y}(t), b_2^{Y}(t), \dots),$$

$$\mathbf{stuckvect}^{Y}(t) = (\mathbf{stuck}_1^{Y}(t), \mathbf{stuck}_2^{Y}(t), \dots).$$

Similarly, for a two-stream externally-jammed process $T = T(Y^A, Y^B)$ we define

$$\mathbf{ballvect}^{T}(t) = (b_1^{T}(t), b_2^{T}(t), \dots),$$

$$\mathbf{stuckvect}^{T}(C, t) = (\mathbf{stuck}_1^{T}(C, t), \mathbf{stuck}_2^{T}(C, t), \dots) \text{ for each } C \in \{A, B\}.$$

We define stuckvect^T(t) as the position-wise sum of stuckvect^T(A, t) and stuckvect^T(B, t).

Definition 24. Suppose that $\mathbf{x} = (x_1, x_2, ...)$ is a sequence of non-negative integers. The noise of \mathbf{x} is defined by $f(\mathbf{x}) = \lambda + \sum_{i>1} x_i p_i$.

Note that for an externally-jammed process Y with arrival rate λ , $f(\mathbf{stuckvect}^Y(t-1))$ is the expected size of $\mathbf{stucksend}^Y(t)$ conditioned on the filtration \mathcal{F}_{t-1}^Y . Similarly, for a two-stream externally-jammed process T with arrival rate λ , $f(\mathbf{stuckvect}^T(t-1))$ is the expected size of $\mathbf{stucksend}^T(A,t) \cup \mathbf{stucksend}^T(B,t)$ conditioned on the filtration \mathcal{F}_{t-1}^T .

3.4 Statement of our technical result

Recall the definition of (λ, η, ν) -suitable (Definition 14). The following Theorem and Corollary form the centre of our argument. Corollary 26 immediately implies Theorem 15, the main technical theorem from the introduction. Theorem 25 and Corollary 26 are proved in Section 6.

Theorem 25. Fix λ , η and ν in (0, 1). Let **p** be a (λ, η, ν) -suitable send sequence. Let T be a two-stream externally-jammed process with arrival rate λ and send sequence **p**. Let $\mathcal{R}^T = \{t \mid \mathbf{stuck}^T(t) = \emptyset\}$. With probability 1, \mathcal{R}^T is finite.

Corollary 26. Fix λ , η and ν in (0, 1). Let **p** be a (λ, η, ν) -suitable send sequence. Let X be a backoff process with arrival rate λ and send sequence **p**. Let $\mathcal{R}^X = \{t \mid \mathbf{balls}^X(t) = \emptyset\}$. Then, with probability 1, \mathcal{R}^X is finite. Hence, X is transient so it is unstable.

3.5 Definitions depending on a (λ, η, ν) -suitable send sequence p

We will need lots of definitions to prove Theorem 25. For everything leading up to and including the proof of Theorem 25, fix an arrival rate $\lambda \in (0, 1)$ a real number $\eta \in (0, 1)$, a real number $\nu \in (0, 1)$, and a (λ, η, ν) -suitable send sequence **p**. Since λ , η and ν are fixed, we will use κ as shorthand for $\kappa(\eta)$ and we will use p_* as a shorthand for $p_*(\lambda, \eta, \nu)$. All of the following notation will depend on **p**, λ , η , and ν .

Definition 27. j_0 is a constant such that for every $j \ge j_0$ we have $p_j > \nu^j$.

Definition 28. $j_{\min} = \min\{j \mid p_j < 1\}.$

We know that j_0 and j_{\min} exist since **p** is (λ, η, ν) -suitable.

Before proceeding with the rest of our formal definitions, we give a brief sense of what their purpose will be. Our proof will proceed by partitioning the bins of T into contiguous "blocks" B_1, B_2, \ldots of exponentially growing length, with block B_i containing roughly κ^i bins. We will wait until, by chance, at some time t_0 , some large initial segment B_1, \ldots, B_{I_0} of bins contains many stuck A-balls and B-balls; we will denote the number of balls required by C_{init} . Since we will choose C_{init} to be large, these stuck balls are likely to take a long time to move out of B_1, \ldots, B_{I_0} , say until time $t_0 + \tau_{\text{init}}$. Until they do so, the *A*-balls will interfere with any balls escaping from Y^B , and crucially they will do so independently from the evolution of Y^B ; likewise, the *B*-balls will interfere with any balls escaping from Y^A .

We will use this interference to argue that B_{I_0+1} is likely to fill with newborn stuck A-balls and B-balls by some time $t_0 + \tau_1 < \tau + \tau_{\text{init}}$, that B_{I_0+2} is likely to fill with stuck A-balls and B-balls by some time $t_0 + \tau_2 < \tau + \tau_{\text{init}}$, and so on. After time $t + \tau_{\text{init}}$, we will show that the process becomes self-sustaining, and that B_i being full of stuck A-balls and B-balls is enough to fill B_{i+1} by itself. We will formalise this idea at the end of the section with the definition of "jammedness".

We now formalise this block construction and define our constants.

Definition 29. For each positive integer *i*, the block B_i is a contiguous interval of positive integers (called bins). This block has bins $\ell(i), \ldots, u(i)$ which are defined as follows. First, u(0) = 0 and $\ell(1) = u(1) = 1$. Then for i > 1, $\ell(i) = u(i - 1) + 1$ and $u(i) = \kappa u(i - 1)$.

Definition 30. The weight of a bin j is $w_j = 1/p_j$. The weight of a block B_i is $W_i = \sum_{j \in B_i} 1/p_j$.

Observe that the weight of a block B_i is the expected number of time steps required for a ball to pass through it, from $b_{\ell(i)}$ to $b_{u(i)+1}$.

Definition 31. $\zeta = \eta \lambda / 24$.

The constant ζ will control how full we require the bins of a given block to be in order to consider T "jammed".

Definition 32. Let I_0 be the smallest integer such that

- $I_0 \geq j_{\min}$.
- For all $n \ge \ell(I_0)$:
 - $|\{j \in [n] \mid p_j \le p_*(\lambda, \eta, \nu)\}| > \eta n$, and - $\nu^n < p_n$.
- $\zeta |B_{I_0}| \ge 4.$
- For $c = \zeta \frac{\kappa 1}{16\kappa^2}$, we have $I_0 \ge \log(4/c)/c$ and $\exp(I_0 c) \ge 4I_0$.

Definition 14 ensures that the second condition can be satisfied. The final condition is satisfied if I_0 is sufficiently large. It ensures that, for all $i \ge I_0$, $4i \le \exp(\zeta i \frac{\kappa-1}{16\kappa^2})$. It is not actually important that I_0 is the smallest integer satisfying the properties in Definition 32, it is merely important that all of these properties (which are monotonically increasing in I_0) are satisfied.

Definition 33. Define $\tau_0 = 0$. For every positive integer *i*, define $\tau_i = \kappa \sum_{k=1}^{I_0+i} (I_0 + i - k + 1) \lceil W_k \rceil$. **Definition 34.** Define $I(0) = I_0 + 1$. For every positive integer τ , $I(\tau) \ge I_0 + 1$ is the integer such that

$$\tau_{I(\tau)-I_0-1} \le \tau < \tau_{I(\tau)-I_0}.$$

Definition 35. For every non-negative integer τ , $bins(\tau) = B_{I(\tau)-1}$.

Intuitively, referring back to the above sketch, $I(\tau)$ is chosen so that at time $t_0 + \tau$, we expect $B_{I(\tau)-1} = bins(\tau)$ to be full of stuck balls and $B_{I(\tau)}$ to be in the process of filling. (Recall that W_i is the expected time required for a ball to travel through block B_i , and observe that $\tau_i - \tau_{i-1} = \kappa \sum_{k=1}^{I_0+i} \lceil W_k \rceil$.)

Definition 36. Let τ_{init} be the smallest integer such that

- $\tau_{\text{init}} \ge \max\{10^7/\lambda^2, 20, (2\kappa/(1-\nu))^4\}.$
- $I(\tau_{\text{init}}) \ge \max\{I_0 + 3, 2I_0(2Q+1)\}$ where $Q = \max_{k=1}^{I_0} \lceil W_k \rceil$.

Once again, it is not important that τ_{init} is the smallest integer satisfying the properties in Definition 36, it is just important that these properties are satisfied, and they can all be satisfied by making τ_{init} sufficiently large, since $I(\tau)$ is $\omega(1)$ as a function of τ .

Definition 37. Let

$$C_{\text{init}} = \left\lceil \max\left(\frac{\zeta |B_{I_0}|}{p_{j_{\min}}}, \frac{12\log(100\tau_{\text{init}})}{(1-p_{j_{\min}})^{\tau_{\text{init}}}}, \frac{2\zeta |\text{bins}(\tau_{\text{init}})|}{p_{j_{\min}}(1-p_{j_{\min}})^{\tau_{\text{init}}}}\right) \right\rceil$$

Recall that $p_{j_{\min}} < 1$ by the definition of j_{\min} , so the definition of C_{\min} in Definition 37 is well-defined.

The following definitions refer to a two-stream externally-jammed process $T = T(Y^A, Y^B)$ with arrival rate λ and send sequence **p**.

Definition 38. For all integers $t \ge 1$, $\mathcal{E}_{init}^T(t)$ is the event that for all $C \in \{A, B\}$, $|\mathbf{stuck}_{j_{\min}}^T(C, t)| \ge C_{init}$.

Definition 39. For $C \in \{A, B\}$ and positive integers t and τ , the two-stream externally-jammed process T is (C, t)-jammed for τ if $f(\mathbf{stuckvect}^T(C, t + \tau - 1)) \ge \zeta |\operatorname{bins}(\tau - 1)|$. It is t-jammed for τ if it is (A, t)-jammed and (B, t)-jammed at τ .

Definition 40. For $C \in \{A, B\}$ and positive integers t and τ , $\mathcal{E}_{jam}^T(C, t, \tau)$ is the event that, for all $\tau' \in \{1, \ldots, \tau\}$, T is (C, t)-jammed for τ' . $\mathcal{E}_{jam}^T(t, \tau)$ is $\mathcal{E}_{jam}^T(A, t, \tau) \cap \mathcal{E}_{jam}^T(B, t, \tau)$.

The heart of our proof will be Lemmas 56, which says that if $\mathcal{E}_{init}^T(t_0)$ occurs, then T is likely to be t_0 -jammed for all $\tau \ge 1$ (and hence never return to a state with no stuck balls). We will split this proof into two regimes. The first regime, $\tau \le \tau_{init}$, will be relatively simple and is covered by Lemma 41; we are essentially just arguing that some of the stuck balls guaranteed by $\mathcal{E}_{init}^T(t_0)$ are likely not to have sent yet. The second regime, $\tau \ge \tau_{init}$, will be the heart of the proof and is covered by Lemma 55.

3.6 The $\tau \leq \tau_{\text{init}}$ regime: Proving Lemma 41

Lemma 41. Fix λ , η and ν in (0,1) and a (λ, η, ν) -suitable send sequence **p**. Let T be a two-stream externally-jammed process with arrival rate λ and send sequence **p**. Consider any $C \in \{A, B\}$. Let t_0 be any positive integer and let $F_{t_0}^T$ be a value of $\mathcal{F}_{t_0}^T$ such that $\mathcal{E}_{init}^T(t_0)$ occurs. Then the following two statements hold.

If
$$\mathcal{F}_{t_0}^T = F_{t_0}^T$$
 occurs then $\mathcal{E}_{jam}^T(C, t_0, 1)$ occurs, and (3)

and

$$\mathbb{P}\left(\mathcal{E}_{jam}^{T}(t_{0},\tau_{init}) \mid \mathcal{F}_{t_{0}}^{T}=F_{t_{0}}^{T}\right) \geq 49/50.$$

$$\tag{4}$$

Proof. We first establish Equation (3). We wish to show that $\mathcal{E}_{jam}^T(C, t_0, 1)$ occurs which is the event that T is (C, t_0) -jammed for time 1, i.e. that $f(\mathbf{stuckvect}^T(C, t_0)) \ge \zeta |\operatorname{bins}(0)| = \zeta |B_{I(0)-1}| = \zeta |B_{I_0}|$. To establish this, first note by the definition of f (Definition 24) that

$$f(\mathbf{stuckvect}^T(C, t_0)) \ge |\mathbf{stuck}_{j_{\min}}^T(C, t_0)|p_{j_{\min}}$$

This quantity can be deduced from $F_{t_0}^T$ and, since $F_{t_0}^T$ guarantees that $\mathcal{E}_{init}^T(t_0)$ happens, is at least $C_{init}p_{j_{\min}}$. To finish, we just need to show that $C_{init}p_{j_{\min}} \geq \zeta |B_{I_0}|$, which follows from the choice of C_{init} in Definition 37.

We next establish Equation (4). Consider any integer τ in the range $1 \leq \tau \leq \tau_{\text{init}}$ and condition on $\mathcal{F}_{t_0}^T = F_{t_0}^T$. Since $\mathcal{E}_{\text{init}}^T(t_0)$ occurs under this conditioning, for each $C \in \{A, B\}$, we have $|\operatorname{stuck}_{j_{\min}}^T(C, t_0)| \geq C_{\text{init}}$; we will argue that $\operatorname{stuck}_{j_{\min}}^T(C, t_0)$ does not have time to empty before time $t_0 + \tau$, so T is likely to be (C, t_0) -jammed for τ . As before, observe that

$$f(\mathbf{stuckvect}^{T}(C, t_{0} + \tau - 1)) \geq p_{j_{\min}} |\mathbf{stuck}_{j_{\min}}^{T}(C, t_{0} + \tau - 1)|$$
$$\geq p_{j_{\min}} |\mathbf{stuck}_{j_{\min}}^{T}(C, t_{0}) \cap \mathbf{stuck}_{j_{\min}}^{T}(C, t_{0} + \tau - 1)|.$$
(5)

For convenience, write $N_{\tau} = |\mathbf{stuck}_{j_{\min}}^T(C, t_0) \cap \mathbf{stuck}_{j_{\min}}^T(C, t_0 + \tau - 1)|$ for the size of the set in the right-hand side of (5).

A ball is in N_{τ} if it is in $\operatorname{stuck}_{j_{\min}}^{T}(C, t_{0})$ and it does not send until time $t_{0} + \tau$ or later. Since messages send independently, conditioned on $\mathcal{F}_{t_{0}}^{T} = F_{t_{0}}^{T}$, N_{τ} is binomially distributed. Let $\mu = \mathbb{E}(|N_{\tau}| | \mathcal{F}_{t_{0}}^{T} = F_{t_{0}}^{T})$, and observe that $\mu = (1 - p_{j_{\min}})^{\tau-1} |\operatorname{stuck}_{j_{\min}}^{T}(C, t_{0})| \geq (1 - p_{j_{\min}})^{\tau_{\min}} C_{\min}$. By the definition of C_{\min} (definition 37), we have $\mu \geq 12 \log(100\tau_{\min})$ and $\mu \geq 2\zeta |\operatorname{bins}(\tau - 1)|/p_{j_{\min}}$. Using this fact, we apply the Chernoff bound of Lemma 19 to $|N_{\tau}|$ with $\delta = 1/2$ to obtain

$$\mathbb{P}(|N_{\tau}| \geq \zeta | \text{bins}(\tau - 1)| / p_{j_{\min}} | \mathcal{F}_{t_0}^T = F_{t_0}^T) \geq \mathbb{P}(|N_{\tau}| \geq \mu/2 | \mathcal{F}_{t_0}^T = F_{t_0}^T) \\ \geq 1 - e^{-\mu/8} \geq 1 - 1/(100\tau_{\text{init}}).$$

Hence by (5) and the definition of (C, t_0) -jammedness, we have that for all $1 \le \tau \le \tau_{\text{init}}$,

$$\mathbb{P}(T \text{ is } (C, t_0) \text{-jammed for time } \tau \mid \mathcal{F}_{t_0}^T = F_{t_0}^T) \geq 1 - 1/(100\tau_{\text{init}}).$$

Thus (4), and therefore also the result, follows by a union bound over all $C \in \{A, B\}$ and $\tau \in \{1, \dots, \tau_{\text{init}}\}$.

3.7 Properties of the block construction

Observation 42. For $i \ge 1$ we have $u(i) = \kappa^{i-1}$. Also $|B_1| = 1$. For $i \ge 2$ we have $|B_i| = \kappa^{i-2}(\kappa - 1) = u(i)(\kappa - 1)/\kappa$.

Proof. For $i \ge 2$, we have $|B_i| = u(i) - u(i-1) = \kappa^{i-1} - \kappa^{i-2}$.

Lemma 43. For all $\tau \ge \tau_{\text{init}}$, $|\text{bins}(\tau)| > \log(\tau)/(2\kappa^2 \log(1/\nu))$.

Proof. By definition, we have $|bins(\tau)| = |B_{I(\tau)-1}|$. By Observation 42 and since $\kappa \ge 2$, it follows that $|bins(\tau)| \ge \kappa^{I(\tau)-3}$. Again by Observation 42, we have $u(I(\tau)) = \kappa^{I(\tau)-1}$, and so

$$|\operatorname{bins}(\tau)| \ge u(I(\tau))/\kappa^2.$$
 (6)

Our proof strategy will be to bound $I(\tau)$ below in order to apply (6). To do this, we will first bound τ above in terms of $u(I(\tau))$, and then convert this into a lower bound on $u(I(\tau))$ in terms of τ . By the definitions of $I(\tau)$ and τ_i , we have

$$\tau < \tau_{I(\tau)-I_0} = \kappa \sum_{k=1}^{I_0} (I(\tau) - k + 1) \lceil W_k \rceil + \kappa \sum_{k=I_0+1}^{I(\tau)} (I(\tau) - k + 1) \lceil W_k \rceil$$

We will argue that the second sum is larger than the first. To see this, let $Q = \max_{k=1}^{I_0} \lceil W_k \rceil$. Then the first sum is at most $\kappa I_0 I(\tau) Q$. To lower-bound the second sum, just consider k up as high as $k = (I(\tau)/2) + 1$ then each term (for each value of k) is at least $\kappa I(\tau)/2$ and the number of terms is at least $(I(\tau)/2) - I_0$. The definition of τ_{init} insures that $I(\tau_{\text{init}}) \geq 2I_0(2Q+1)$. Thus, for every $\tau \geq \tau_{\text{init}}$, the second term is at least as large as the first.

It follows that

$$\tau < 2\kappa \sum_{k=I_0+1}^{I(\tau)} (I(\tau) - k + 1) \lceil W_k \rceil = 2\kappa \sum_{k=I_0+1}^{I(\tau)} \sum_{m=I_0+1}^k \lceil W_m \rceil = 2\kappa \sum_{k=I_0+1}^{I(\tau)} \left\lceil \sum_{j=\ell(I_0+1)}^{u(k)} (1/p_j) \right\rceil.$$

We bound the right-hand side above by observing that

$$\{u(k): I_0 + 1 \le k \le I(\tau)\} \subseteq \{\ell(I_0 + 1), \ell(I_0 + 1) + 1, \dots, u(I(\tau))\},\$$

and so

$$\tau < 2\kappa \sum_{a=\ell(I_0+1)}^{u(I(\tau))} \left[\sum_{j=\ell(I_0+1)}^{a} (1/p_j) \right] \le 4\kappa \sum_{a=\ell(I_0+1)}^{u(I(\tau))} \sum_{j=\ell(I_0+1)}^{a} (1/p_j).$$
(7)

By the definition of I_0 , for all $i \ge I_0 + 1$ and all $j \in B_i$ we have $p_j > \nu^j$. It follows that for all $a \ge \ell(I_0 + 1)$,

$$\sum_{j=\ell(I_0+1)}^{a} \frac{1}{p_j} < \sum_{j=\ell(I_0+1)}^{a} \left(\frac{1}{\nu}\right)^j < \left(\frac{1}{\nu}\right)^a \sum_{j=0}^{\infty} \nu^j = \frac{1}{1-\nu} \left(\frac{1}{\nu}\right)^a.$$

Hence by (7),

$$\tau < \frac{4\kappa}{1-\nu} \sum_{a=\ell(I_0+1)}^{u(I(\tau))} \left(\frac{1}{\nu}\right)^a < 4\kappa \frac{\left(\frac{1}{\nu}\right)^{u(I(\tau))}}{1-\nu} \sum_{a=0}^{\infty} \nu^a = \frac{4\kappa}{\left(1-\nu\right)^2} \left(\frac{1}{\nu}\right)^{u(I(\tau))}.$$
(8)

Solving for $u(I(\tau))$ now yields

$$u(I(\tau)) > \log_{1/\nu}(\tau) + 2\log_{1/\nu}\left(\frac{1-\nu}{2\sqrt{\kappa}}\right).$$

Since $\tau \ge \tau_{\text{init}}$ and $\tau_{\text{init}} \ge (2\kappa/(1-\nu))^4$, it follows that $u(I(\tau)) > \frac{1}{2} \log_{1/\nu} \tau$. By (6), it follows that

$$|\operatorname{bins}(\tau)| \ge \frac{1}{\kappa^2} u(I(\tau)) > \frac{1}{2\kappa^2 \log(1/\nu)} \log \tau,$$

as required.

3.8 Noise

Lemma 44. Let $T(Y^A, Y^B)$ be a two-stream externally-jammed process with arrival rate λ and send sequence **p** with $p_0 = 1$. Let $t \ge 0$ and let $C \in \{A, B\}$. Let **x** be any possible value of $\mathbf{stuckvect}^T(C, t)$. Then

$$\mathbb{P}(\mathbf{stucksend}^T(C, t+1) = \emptyset \mid \mathbf{stuckvect}^T(C, t) = \mathbf{x}) \le \exp(-f(\mathbf{x})/3).$$

Proof. For brevity, write q for the probability in the statement of the lemma, and let Ψ be the number of balls in $\mathbf{stuck}^T(C, t) \cap \mathbf{send}^T(t+1)$. Then since any ball in $\mathbf{stucksend}^T(C, t+1)$ is either born at time t+1 or in Ψ , we have

$$q = e^{-\lambda/2} \mathbb{P}(\Psi = 0 \mid \mathbf{stuckvect}^T(C, t) = \mathbf{x}).$$
(9)

Conditioned on $\mathbf{stuckvect}^T(C,t) = \mathbf{x}$, Ψ is a sum of independent Bernoulli variables with mean $\mu = f(\mathbf{x}) - \lambda$. If one of these Bernoulli variables has parameter p then the probability that it is 0 is $1 - p \leq \exp(-p)$. Thus the probability that they are all 0 is at most $\exp(-\mu)$. We therefore have

$$\mathbb{P}(\Psi = 0 \mid \mathbf{stuckvect}^T(C, t) = \mathbf{x}) \le e^{-\mu}.$$

It follows by (9) that $q \le e^{-f(\mathbf{x}) + \lambda/2} < e^{-f(\mathbf{x})/3}$ as required (since $f(\mathbf{x}) \ge \lambda$ by definition).

4 Towards a proof of Theorem 25: Jamming causes a product distribution of stuck balls in a two-stream externally-jammed process

The purpose of this section is to prove the following lemma, which will be our main tool to analyse the two-stream process in the proof of Lemma 55 (and thereby Theorem 25).

Lemma 45. Fix λ , η and ν in (0,1) and a (λ, η, ν) -suitable send sequence \mathbf{p} with $p_0 = 1$. Let Y^A and Y^B be independent externally-jammed processes with arrival rate $\lambda/2$ and send sequence \mathbf{p} and consider the two-stream externally-jammed process $T = T(Y^A, Y^B)$. Consider any integer $t_0 \geq j_{\min}$. Let $F_{t_0}^T$ be a value of $\mathcal{F}_{t_0}^T$ such that $\mathcal{E}_{\min}^T(t_0)$ happens. Consider any integer $\tau \geq \tau_{\min}$. Then there is a coupling of

- T conditioned on $\mathcal{F}_{t_0}^T = F_{t_0}^T$,
- a sample $\{Z_j^A \mid j \in bins(\tau)\}$ where each Z_j^A is chosen independently from a Poisson distribution with mean $\lambda/(4p_j)$, and
- a sample $\{Z_j^B \mid j \in bins(\tau)\}$ where each Z_j^B is chosen from a Poisson distribution with mean $\lambda/(4p_j)$ and these are independent of each other but not of the $\{Z_j^A\}$ values

in such a way that at least one of the following happens:

- $\mathcal{E}_{iam}^T(t_0, \tau)$ does not occur, or
- for all $j \in bins(\tau)$, $|\mathbf{stuck}_j^T(A, t_0 + \tau)| \ge Z_j^A$ and $|\mathbf{stuck}_j^T(B, t_0 + \tau)| \ge Z_j^B$.

In order to prove Lemma 45 we need to study several stochastic processes, starting with a randomunsticking process in the next section.

4.1 The definition of a random-unsticking process

Let Y be an externally-jammed process with arrival rate λ and send sequence **p**. Let t_0 be a fixed positive integer. We will define a *random-unsticking process* $R = R(Y, t_0)$ by redefining the "stuck" and "unstuck" states of Y, so that balls in R become unstuck independently at random with probability depending on t_0 . Later, in Lemma 51, we will show that R and T can be coupled in such a way that as long as T stays jammed, if a ball is stuck in R then it is also stuck in T; this will allow us to analyse R rather than T in proving Lemma 45.

Apart from this redefinition, which explained below, everything about R is the same as Y: for all $t \ge 0$, $\mathbf{balls}^R(t) = \mathbf{balls}^T(t)$; for all $t \ge 0$ and $j \ge 1$, $b_j^R(t) = b_j^Y(t)$; for all $t \ge 1$ and $j \ge 0$, $b_j'^R(t) = b_j'^Y(t)$; for all $t \ge 1$ send^R $(t) = \mathbf{send}^Y(t)$; and for $t \ge 0$ and all $\beta \in \mathbf{balls}^T(t)$, $\mathbf{birth}^R(\beta) = \mathbf{birth}^Y(\beta)$.

We now formalise the new "stuck" and "unstuck" states in R. The set $\mathbf{stuck}^R(t)$ will be the set of all balls in $\mathbf{balls}^R(t)$ that are stuck in the process R at time t and $\mathbf{stuck}_j^R(t) = \mathbf{stuck}^R(t) \cap b_j^T(t)$. As the reader will expect, $\mathbf{unstuck}^R(t) = \mathbf{balls}^R(t) \setminus \mathbf{stuck}^R(t)$ so for each t we only need to define $\mathbf{stuck}^R(t)$ or $\mathbf{unstuck}^R(t)$ and the other is defined implicitly. These variables are defined as follows. First, let $\mathbf{stuck}^R(0) = \emptyset$ and

$$p_{\text{unstick}}(t) = \begin{cases} 1, & \text{if } t \le t_0 \\ e^{-\zeta |\text{bins}(t-t_0)|/16}, & \text{otherwise.} \end{cases}$$

Then for any $t \ge 1$ we have the following definitions:

- **newstuck**^R(t) =**stuck**^R $(t-1) \cup b_0^{\prime R}(t)$.
- stucksend^R(t) =newstuck^R $(t) \cap$ send^R(t).
- For each β ∈ balls^R(t), unstick(β, t) is an independent Bernoulli random variable which is 1 with probability p_{unstick}(t) and 0 otherwise.
- Then $\mathbf{unstuck}^{R}(t) = \mathbf{unstuck}^{R}(t-1) \cup \{\beta \in \mathbf{stucksend}^{R}(t) \mid \mathbf{unstick}(\beta, t) = 1\}.$

We take \mathcal{F}_t^R to be \mathcal{F}_t^Y along with the values of the random variables $\mathbf{unstick}(\beta, t')$ for $t' \leq t$ and $\beta \in \mathbf{balls}^R(t')$ so that all information about the evolution of the process R, including the value of $\mathbf{stuck}^R(t)$, is captured.

4.2 Defining a reverse random-unsticking process

Given an externally-jammed process Y with arrival rate λ , a positive integer t_0 , a random-unsticking process $R = R(Y, t_0)$, and positive integers j and τ_{end} , we wish to find a lower bound on $|\mathbf{stuck}_j^R(t_0 + \tau_{end})|$ in order to prove Lemma 45. To do this, we'll define the notion of a reverse random-unsticking process, which is essentially a time-reversal of a random-unsticking process; we will make this relation formal later, in Lemma 47. The reverse random-unsticking process $\widetilde{R} = \widetilde{R}(Y, t_0, \tau_{end})$ will be easier to study than R itself, allowing us to deduce the lower bounds on $|\mathbf{stuck}_j^R(t_0 + \tau_{end})|$ that we need. Note that the basic idea of using a time-reversal in the context of backoff processes is due to Aldous [4]. Here we require independent lower bounds for different j, so it is a lot more complicated.

The process R will run for steps $1, \ldots, \tau_{\text{end}}$. Each step consists of (II) sending, (III) adjusting the bins, and (IV) random unsticking. (We start counting from (II) so that the indices line up with the parts of the steps of an externally-jammed process. Here there is no step initialisation, since there are no births. On the other hand, there is random unsticking.) Let $J_{\text{max}} = u(I(\tau_{\text{end}}) - 1) + \tau_{\text{end}}$. The bins in the process \tilde{R} will be the integers $j \in [J_{\text{max}}]$. For $j \in [J_{\text{max}}]$, the set $b_j^{\tilde{R}}(\tau)$ is the population of bin j just after (all parts of) the τ 'th step. For convenience, we define $b_{J_{\text{max}}+1}^{\tilde{R}}(\tau) = \emptyset$. Also, $\text{balls}^{\tilde{R}}(\tau) = \bigcup_{j=1}^{J_{\text{max}}} b_j^{\tilde{R}}(\tau)$.

The process \widetilde{R} is initialised as follows. For each $j \in [J_{\max}]$, the number \widetilde{x}_j of balls that start in $b_j^{\widetilde{R}}(0)$ is chosen independently from a Poisson distribution with mean λ/p_j . These balls are given unique names $(\widetilde{R}, j, 1), \ldots, (\widetilde{R}, j, \widetilde{x}_j)$. Step τ is defined as follows.

Part (II) of step τ: For j ∈ [J_{max}], all balls in b_j^R(τ − 1) send independently with probability p_j. We use send^R(τ) for the set of balls that send at time τ.

• Part (III) of step τ : For all $j \in [J_{\max}]$,

$$b_{j}^{\widetilde{R}}(\tau) = (b_{j+1}^{\widetilde{R}}(\tau-1) \cap \mathbf{send}^{\widetilde{R}}(\tau)) \cup (b_{j}^{\widetilde{R}}(\tau-1) \setminus \mathbf{send}^{\widetilde{R}}(\tau)).$$

Part (IV) of step τ: For each β ∈ balls^{*˜*}(τ − 1), unstick(β, τ) is an independent Bernoulli random variable, which is 1 with probability p_{unstick}(t₀ + τ_{end} − τ + 1) and 0 otherwise.

Note that balls in $b_1^{\widetilde{R}}(\tau - 1)$ that send at step τ are not part of $\mathbf{balls}^{\widetilde{R}}(\tau)$ — instead, they leave the system at step τ . Note also that we do not define "stuck" and "unstuck" states for \widetilde{R} — it will actually be more convenient to work with the **unstick** variables directly. (These will later be coupled to the corresponding variables in a random-unsticking process \widetilde{R} ; as the definition suggests, unsticking events at times $1, 2, \ldots, \tau_{\text{end}}$ in \widetilde{R} will correspond to unsticking events at times $t_0 + \tau_{\text{end}}, t_0 + \tau_{\text{end}} - 1, \ldots, t_0 + 1$ in R.)

4.3 Defining trajectories in a random-unsticking process

Let Y be an externally-jammed process with arrival rate λ and let t_0 be a positive integer. Consider a random-unsticking process $R = R(Y, t_0)$.

We will only be concerned with balls β with birth^{*R*}(β) $\geq t_0 + 1$. For $t \geq t_0 + 1$ consider a ball with name $\beta = (Y, t, x)$ for some x. The *ball trajectory* B of the ball β (up to time step $t_0 + \tau_{end}$ for a positive integer τ_{end}) contains the following information:

- $t_{\text{birth}}(B) = \text{birth}^R(\beta).$
- The bin J(B) such that $\beta \in b_{J(B)}^R(t_0 + \tau_{end})$.
- Positive integers $N_1(B), \ldots, N_{J(B)}(B)$ adding up to $t_0 + \tau_{end} t_{birth}(B) + 1$. For each $j \in [J(B)]$, $N_j(B) = |\{t' \mid t_{birth}(B) \le t' \le t_0 + \tau_{end}, \ \beta \in b_j^R(t')\}|.$
- For each integer t' in the range t_{birth}(B) ≤ t' ≤ t₀ + τ_{end}, the indicator variable unstick(B,t') = unstick(β,t').

Note that any choice of positive integers $N_1(B), \ldots, N_{J(B)}(B)$ with the right sum leads to a valid (ball) trajectory, and so does any sequence of indicator random variables $\operatorname{unstick}(B, t')$. Thus, for all fixed λ , t_0 and $\tau_{\operatorname{end}}$, all of the information in trajectory B is captured by J(B), $t_{\operatorname{birth}}(B)$, $\{N_j(B) \mid 1 \le j \le J(B)\}$ and $\{\operatorname{unstick}(B, t') \mid t_{\operatorname{birth}}(B) \le \tau \le t_0 + \tau_{\operatorname{end}}\}$. Let $\mathcal{B}^R(t_0, \tau_{\operatorname{end}}, t, J)$ be the set of trajectories up to $t_0 + \tau_{\operatorname{end}}$ with $t_{\operatorname{birth}}(B) = t$ and J(B) = j.

Given a trajectory B, let $S(B) = \{t_{\text{birth}}(B) + \sum_{k=1}^{c} N_k(B) \mid c \in \{0, \dots, J(B) - 1\}\}$. We will use the fact that S(B) is the set of time steps up to time $t_0 + \tau_{\text{end}}$ at which balls following trajectory B send.

We can calculate the distribution of the number of balls following a given trajectory B. This is a Poisson random variable whose mean, $\mu^R(B)$, is given by $\mu^R(B) = F_1F_2F_3$ where F_1 , F_2 and F_3 are defined as follows:

$$F_{1} = \lambda (1 - p_{J(B)})^{N_{J(B)}(B)-1}; \qquad F_{2} = \prod_{j=1}^{J(B)-1} p_{j} (1 - p_{j})^{N_{j}(B)-1};$$

$$F_{3} = \prod_{t'=t_{\text{birth}}(B)}^{t_{0}+\tau_{\text{end}}} p_{\text{unstick}}(t')^{\text{unstick}(B,t')} (1 - p_{\text{unstick}}(t'))^{1-\text{unstick}(B,t')}.$$

4.4 Defining trajectories in a reverse random-unsticking process R

Let Y be an externally-jammed process with arrival rate λ , let t_0 be a positive integer and let τ_{end} be a positive integer. Consider a reverse random-unsticking process $\widetilde{R} = \widetilde{R}(Y, t_0, \tau_{\text{end}})$

We will only be concerned with balls in \overline{R} that leave the process by sending from bin 1 at some step $\tau \in \{1, \ldots, \tau_{\text{end}}\}$. We next define the trajectory \widetilde{B} of such a ball with name $\beta = (\widetilde{R}, j', x)$ in \widetilde{R} . The trajectory \widetilde{B} contains the following information:

- The bin $J(\widetilde{B}) = j'$, the bin into which β is placed in the initialisation, so $\beta \in b_{J(\widetilde{B})}^{\widetilde{R}}(0)$.
- The time $\tau_{\text{leave}}(\widetilde{B})$ at which β leaves the process by sending from bin 1, so $\beta \in b_1^{\widetilde{R}}(\tau_{\text{leave}}(\widetilde{B}) 1)$ but $\beta \notin b_1^{\widetilde{R}}(\tau_{\text{leave}}(\widetilde{B})).$
- Positive integers $N_1(\widetilde{B}), \ldots, N_{J(\widetilde{B})}(\widetilde{B})$, adding up to $\tau_{\text{leave}}(\widetilde{B})$. For each $j \in [J(\widetilde{B})]$, $N_j(\widetilde{B}) = |\{\tau' \mid 0 \le \tau' \le \tau_{\text{leave}}(\widetilde{B}) 1, \ \beta \in b_j^{\widetilde{R}}(\tau')\}|$.
- For each time step τ' in the range 1 ≤ τ' ≤ τ_{leave}(B̃), the indicator variable unstick(B̃, τ') = unstick(β, τ').

Note that any choice of positive integers $N_1(\widetilde{B}), \ldots, N_{J(\widetilde{B})}(\widetilde{B})$ with the right sum leads to a valid trajectory, and so does any sequence of indicator random variables $\operatorname{unstick}(\widetilde{B}, \tau)$. Thus for all fixed λ , t_0 and $\tau_{\operatorname{end}}$, all of the information in trajectory B is captured by $J(\widetilde{B}), \tau_{\operatorname{leave}}(\widetilde{B}), \{N_j(\widetilde{B}) \mid 1 \leq j \leq J(\widetilde{B})\}$ and $\{\operatorname{unstick}(\widetilde{B}, \tau) \mid 1 \leq \tau \leq \tau_{\operatorname{leave}}(\widetilde{B})\}$. Let $\mathcal{B}^{\widetilde{R}}(t_0, \tau_{\operatorname{end}}, \tau, J)$ be the set of trajectories with $\tau_{\operatorname{leave}}(\widetilde{B}) = \tau$ and $J(\widetilde{B}) = J$.

Given a trajectory \widetilde{B} , let $\widetilde{S}(\widetilde{B}) = \{\sum_{k=\widetilde{c}}^{J(\widetilde{B})} N_k(\widetilde{B}) \mid \widetilde{c} \in \{1, \ldots, J(\widetilde{B})\}\}$. We will use the fact that $\widetilde{S}(\widetilde{B})$ is the set of time steps at which balls following trajectory \widetilde{B} send.

In the process \hat{R} , the distribution for the number of balls following trajectory \tilde{B} is a Poisson random variable whose mean, $\mu^{\tilde{R}}(\tilde{B})$, is given by $\mu^{\tilde{R}}(\tilde{B}) = \tilde{F}_1 \tilde{F}_2 \tilde{F}_3$ where \tilde{F}_1 , \tilde{F}_2 and \tilde{F}_3 are defined as follows.

$$\tilde{F}_{1} = \frac{\lambda}{p_{J(\tilde{B})}}; \qquad \tilde{F}_{2} = \prod_{j=1}^{J(\tilde{B})} p_{j}(1-p_{j})^{N_{j}(\tilde{B})-1}; \\
\tilde{F}_{3} = \prod_{\tau=1}^{\tau_{\text{leave}}(\tilde{B})} p_{\text{unstick}}(\tau_{\text{end}} - \tau + t_{0} + 1)^{\mathbf{unstick}(\tilde{B},\tau)} (1-p_{\text{unstick}}(\tau_{\text{end}} - \tau + t_{0} + 1))^{1-\mathbf{unstick}(\tilde{B},\tau)}.$$

4.5 A probability-preserving bijection from trajectories of a random-unsticking process R to those of the corresponding reverse random-unsticking process \widetilde{R}

Let Y and Y' be externally-jammed processes with arrival rate λ and send sequence p, and let t_0 and τ_{end} be positive integers. Consider a random-unsticking process $R = R(Y, t_0)$ and a reverse random-unsticking process $\widetilde{R} = \widetilde{R}(Y', t_0, \tau_{end})$ We now set out the natural bijection from trajectories of R to trajectories of \widetilde{R} that will later form the basis of our coupling.

Definition 46. The time-reversal bijection π from $\mathcal{B}^R(t_0, \tau_{end}, t, J)$ to $\mathcal{B}^{\widetilde{R}}(t_0, \tau_{end}, \tau_{end} - t + t_0 + 1, J)$ is defined as follows for all $B \in \mathcal{B}^R(t_0, \tau_{end}, t, J)$:

• $\tau_{\text{leave}}(\pi(B)) = \tau_{\text{end}} - t_{\text{birth}}(B) + t_0 + 1;$

- $J(\pi(B)) = J;$
- for all $k \in [J]$, $N_k(\pi(B)) = N_k(B)$;
- for all t' in the range $t_{\text{birth}}(B) \le t' \le t_0 + \tau_{\text{end}}$, $\mathbf{unstick}(\pi(B), \tau_{\text{end}} t' + t_0 + 1) = \mathbf{unstick}(B, t')$.

We now pause for a moment to observe that π is a valid map into $\mathcal{B}^{\widetilde{R}}(t_0, \tau_{\text{end}}, \tau_{\text{end}} - t + t_0 + 1, J)$. Indeed, we have $t_{\text{birth}}(B) \ge t_0 + 1$, so $\tau_{\text{leave}}(\pi(B)) \le \tau_{\text{end}}$ as required. In order for $J(\pi(B)) = J$ to be a bin of the process \widetilde{R} we need J to be in $[J_{\text{max}}]$; this is fine since the definition of J(B) = J ensures that $J(B) \le \tau_{\text{end}}$ and the definition of J_{max} ensures that $1 \le \tau_{\text{end}} \le J_{\text{max}}$ so $J(B) \le J_{\text{max}}$. We also have from the definitions that

$$\sum_{j=1}^{J} N_j(B) = \tau_{\text{end}} - t_{\text{birth}}(B) + t_0 + 1 = \tau_{\text{leave}}(B) = \sum_{j=1}^{J} N_j(\pi(B)),$$

as required. Finally, the values of τ for which we define $\operatorname{unstick}(\pi(B), \tau)$ range from $\tau_{\operatorname{end}} - t_{\operatorname{birth}}(B) + t_0 + 1 = \tau_{\operatorname{leave}}(\pi(B))$ down to 1, as required. Thus the image of π is indeed contained in $\mathcal{B}^{\widetilde{R}}(t_0, \tau_{\operatorname{end}}, \tau_{\operatorname{end}} - t + t_0 + 1, J)$, and it is clear that π is a bijection as claimed.

The following lemma formalises our intuition that \overline{R} is a time reversal of R, and will later form the basis for our coupling between them.

Lemma 47. Let Y and Y' be externally-jammed processes with arrival rate λ and send sequence **p** and let t_0 and τ_{end} be positive integers. Consider a random-unsticking process $R = R(Y, t_0)$ and a reverse random-unsticking process $\tilde{R} = \tilde{R}(Y', t_0, \tau_{end})$. Let t be an integer satisfying $t \ge t_0 + 1$. Let J be a positive integer. Let π be the time-reversal bijection from $\mathcal{B}^R(t_0, \tau_{end}, t, J)$ to $\mathcal{B}^{\tilde{R}}(t_0, \tau_{end}, -t + t_0 + 1, J)$. Then for any $B \in \mathcal{B}^R(t_0, \tau_{end}, t, J)$, we have $\mu^R(B) = \mu^{\tilde{R}}(\pi(B))$ and $\tilde{S}(\pi(B)) = \{\tau_{end} - t + t_0 + 1 \mid t \in S(B)\}$.

Proof. It is straightforward to use the definitions of π , μ^R and $\mu^{\widetilde{R}}$ to see that $\mu^R(B) = \mu^{\widetilde{R}}(\pi(B))$. The only subtlety is checking that the probabilities of the unsticking indicator variables correspond. For this, note from the definition of π that for $\tau \in \{1, \ldots, \tau_{\text{leave}}(\pi(B))\}$, the probability that $\text{unstick}(\pi(B), \tau) = 1$ is $p_{\text{unstick}}(t')$ where $\tau = \tau_{\text{end}} - t' + t_0 + 1$. So this probability is $p_{\text{unstick}}(\tau_{\text{end}} - \tau + t_0 + 1)$, as required.

To see that $S(\pi(B)) = \{\tau_{end} - t + t_0 + 1 \mid t \in S(B)\}$, note that if we take $t = t_{birth}(B) + \sum_{k=1}^{c} N_k(B)$ in S(B) (for some $c \in \{0, \dots, J(B) - 1\}$) then the quantity $\tau_{end} - t + t_0 + 1$ is equal to $\tau_{end} - t_{birth}(B) - \sum_{k=1}^{c} N_k(B) + t_0 + 1$ which, using the fact that the sum of all of the $N_k(B)$'s is $\tau_{end} - t_{birth}(B) + t_0 + 1$, is

$$\tau_{\text{end}} - t_{\text{birth}}(B) - (\tau_{\text{end}} - t_{\text{birth}}(B) + t_0 + 1) + \sum_{k=c+1}^{J(B)} N_k(t) + t_0 + 1.$$

Clearly, this is $\sum_{k=\tilde{c}}^{J(\tilde{B})} N_k(\tilde{B})$ for $\tilde{c} = c + 1$, which is in $\tilde{S}(\pi(B))$.

4.6 The set of balls that "fill" bin *j*

Let Y be an externally-jammed process with arrival rate λ and let t_0 be a positive integer. Consider a random-unsticking process $R = R(Y, t_0)$. The following definitions will be convenient.

For any positive integer j in block B_i and any positive integer τ , we define $\operatorname{Fill}_j^R(\tau)$ to be the subset of balls $\beta \in \operatorname{stuck}_i^R(t_0 + \tau)$ such that

birth^R(
$$\beta$$
) $\geq \max\left\{t_0 + 1, t_0 + \tau - \kappa \sum_{k=1}^{i} \lceil W_k \rceil\right\}$.

We make a similar definition for the process $T = T(Y^A, Y^B)$. For any $C \in \{A, B\}$, any positive integer j in B_i , and any positive integer τ , $\operatorname{Fill}_j^T(C, \tau)$ is the subset of balls $\beta \in \operatorname{stuck}_j^T(C, t_0 + \tau)$ such that

$$\operatorname{birth}^{T}(\beta) \geq \max\left\{t_{0} + 1, t_{0} + \tau - \kappa \sum_{k=1}^{i} \lceil W_{k} \rceil\right\}$$

It will be convenient to think about $\operatorname{Fill}_{j}^{R}(\tau)$ in terms of trajectories, so we make the following equivalent definition for $\operatorname{Fill}_{j}^{R}(\tau)$: For any positive integer j in block B_{i} and any positive integer τ_{end} , we define $\operatorname{Fill}_{j}^{R}(\tau_{\mathrm{end}})$ to be the set of balls that take trajectories B satisfying the following:

- (i) J(B) = j.
- (ii) $t_{\text{birth}}(B) \ge t_0 + \tau_{\text{end}} \kappa \sum_{k=1}^{i} \lceil W_k \rceil$.
- (iii) $t_{\text{birth}}(B) \ge t_0 + 1$.
- (iv) for all $t' \in S(B)$, unstick(B, t') = 0.

Finally, we define a corresponding notion $\operatorname{Fill}_{j}^{\widetilde{R}}(\tau_{\operatorname{end}})$ for a reverse random-unsticking process. We first define $\operatorname{Fill}_{j}^{\widetilde{R}}(\tau_{\operatorname{end}})$ in terms of trajectories. Let Y be an externally-jammed process with arrival rate λ and let t_{0} and $\tau_{\operatorname{end}}$ be positive integers. Consider a reverse random-unsticking process $\widetilde{R} = \widetilde{R}(Y, t_{0}, \tau_{\operatorname{end}})$. Let j be a positive integer in block B_{i} . Then $\operatorname{Fill}_{j}^{\widetilde{R}}(\tau_{\operatorname{end}})$ is defined to be the set of balls that take trajectories \widetilde{B} satisfying the following.

(i) $J(\widetilde{B}) = j$.

(ii)
$$\tau_{\text{leave}}(\widetilde{B}) \le 1 + \kappa \sum_{k=1}^{i} \lceil W_k \rceil$$
.

- (iii) $\tau_{\text{leave}}(\widetilde{B}) \leq \tau_{\text{end}}.$
- (iv) for all $\tau' \in \widetilde{S}(\widetilde{B})$, $\mathbf{unstick}(\widetilde{B}, \tau') = 0$.

Equivalently, $\operatorname{Fill}_{i}^{\widetilde{R}}(\tau_{\operatorname{end}})$ is the set of balls $\beta = (\widetilde{R}, j, x')$ satisfying the following.

- (I) $\beta \in b_i^{\widetilde{R}}(0)$.
- (II) For some $\tau' \leq 1 + \kappa \sum_{k=1}^{i} \lceil W_k \rceil$, $\beta \in b_1^{\widetilde{R}}(\tau'-1)$ and β sends at time τ' .
- (III) This τ' also satisfies $\tau' \leq \tau_{end}$.
- (IV) β does not send on any step $\tau \in \{1, 2, \dots, 1 + \kappa \sum_{k=1}^{i} \lceil W_k \rceil\}$ with $unstick(\beta, \tau) = 1$.

4.7 Bounding $\operatorname{Fill}_{j}^{R}(\tau_{\operatorname{end}})$

Lemma 48. Let Y be an externally-jammed process with arrival rate λ and send sequence **p**. Consider any integer $t_0 \ge j_{\min}$. Let $\tau_{end} \ge \tau_{init}$ be an integer. Then there is a coupling of

- (i) a reverse random-unsticking process $\widetilde{R}(Y, t_0, \tau_{end})$ with
- (ii) a sample $\{Z_j^- \mid j \in bins(\tau_{end})\}$ where each Z_j^- is chosen independently from a Poisson distribution with mean $\lambda/(4p_j)$

in such a way that, for every integer $j \in B_{I(\tau_{end})-1}$, $|\operatorname{Fill}_{j}^{\widetilde{R}}(\tau_{end})| \geq Z_{j}^{-}$.

Proof. Let $\widetilde{R} = \widetilde{R}(Y, t_0, \tau_{end})$. We will use the definition of $\operatorname{Fill}_{j}^{\widetilde{R}}(\tau_{end})$ without trajectories. To simplify the notation, let $i = I(\tau_{end}) - 1$, so that $\operatorname{bins}(\tau_{end}) = B_i$. Since $\tau_{end} \geq \tau_{init}$, and since the definition of τ_{init} (Definition 36) ensures that $I(\tau_{init}) \geq I_0 + 3$, we have $i \geq I_0 + 2$.

We will be considering $j \in B_i$. Note that this ensures that $j \leq u(i)$ and this is at most J_{\max} , by the definition of J_{\max} in Section 4.2. Thus, by the initialisation of the process \tilde{R} , the number of balls in $b_j^{\tilde{R}}(0)$ is Poisson with mean λ/p_j (independently of any other j). Since each $\operatorname{Fill}_j^{\tilde{R}}(\tau_{\operatorname{end}})$ is determined by the trajectories of balls in $b_j^{\tilde{R}}(0)$, and since all balls behave independently in the process \tilde{R} , the random variables $\operatorname{Fill}_j^{\tilde{R}}(0)$ are already distributed as independent Poisson random variables as in (ii). So to finish, we just need to show that, for all $j \in B_i$ and all balls $\beta \in b_j^{\tilde{R}}(0)$, the probability that (II), (III) and (IV) happens is at least 1/4. To do this, we'll show that the probability that (II) or (III) fails is at most 1/2 and that the probability that (IV) fails is at most 1/4 (then we are finished by a union bound adding the failure probabilities 1/2 and 1/4).

For (IV), it suffices to show that for any $\tau \in \{1, \ldots, 1 + \kappa \sum_{k=1}^{i} \lceil W_k \rceil\}$, $p_{\text{unstick}}(t_0 + \tau_{\text{end}} - \tau + 1) \leq 1/(4j)$. Using a union bound over the (at most j) steps when β sends, this implies that the probability that β ever has $\text{unstick}(\beta, \tau) = 1$ on one of these steps τ is at most 1/4. Towards this end, first note from the definitions of $I(\cdot)$ and $\tau_{i'}$ that

$$\tau_{\text{end}} \ge \tau_{I(\tau_{\text{end}})-I_0-1} = \tau_{i-I_0} = \tau_{i-I_0-1} + \kappa \sum_{k=1}^{i} \lceil W_k \rceil.$$
(10)

In particular, (10) implies that $t_0 + \tau_{end} - \tau + 1 \ge t_0 + 1$ for all $\tau \le 1 + \kappa \sum_{k=1}^{i} \lceil W_k \rceil$. It follows from the definitions of $p_{unstick}(\cdot)$ and $bins(\cdot)$ that

$$p_{\text{unstick}}(t_0 + \tau_{\text{end}} - \tau + 1) = \exp(-\zeta |\text{bins}(\tau_{\text{end}} - \tau + 1)|/16) = \exp(-\zeta |B_{I(\tau_{\text{end}} - \tau + 1)-1}|/16).$$

Since $|B_{i'}|$ is monotonically increasing in i' (which follows from $\kappa > 2$) and $I(\cdot)$ is monotonically increasing in its argument, the right hand side is maximised when τ is as large as possible, so

$$p_{\text{unstick}}(t_0 + \tau_{\text{end}} - \tau + 1) \le \exp(-\zeta |B_{I(\tau_{\text{end}} - \kappa \sum_{k=1}^i \lceil W_k \rceil) - 1}|/16).$$

$$\tag{11}$$

To simplify the notation, let $x = \tau_{\text{end}} - \kappa \sum_{k=1}^{i} \lceil W_k \rceil$. From the definition of $I(\cdot)$ we have $x < \tau_{I(x)-I_0}$, so since $x \ge \tau_{(i-1)-I_0}$ by (10) we have i - 1 < I(x) and hence $i \le I(x)$. Plugging in the value of x, we conclude that $I(\tau_{\text{end}} - \kappa \sum_{k=1}^{i} \lceil W_k \rceil) \ge i$. Plugging this into (11), we get

$$p_{\text{unstick}}(t_0 + \tau_{\text{end}} - \tau + 1) \le \exp(-\zeta |B_{i-1}|/16).$$

Since $i \ge I_0 \ge 4$, Observation 42 gives $|B_{i-1}| \ge u(i-1)(\kappa-1)/\kappa$, and we can simplify this as follows:

$$p_{\text{unstick}}(t_0 + \tau_{\text{end}} - \tau + 1) \le \frac{1}{e^{\zeta u(i-1)\frac{\kappa-1}{16\kappa}}} = \frac{1}{e^{\zeta u(i)\frac{\kappa-1}{16\kappa^2}}}.$$
(12)

Since $u(i) \ge i \ge I_0$, the definition of I_0 ensures that $\exp(\zeta u(i)\frac{\kappa-1}{16\kappa^2}) \ge 4u(i)$; we therefore have $p_{\text{unstick}}(t_0 + \tau_{\text{end}} - \tau + 1) \le 1/(4u(i)) \le 1/(4j)$, as required for the union bound.

For (II) and (III) we first note that (II) implies (III) by (10). So to finish the proof, we will prove the probability that (II) fails for a given ball $\beta \in b_j^{\tilde{R}}(0)$ is at most 1/2. The number of time steps that it takes for β to send from bin 1 is a random variable Ψ which is the sum of independent geometric random variables

with parameters p_j, \ldots, p_1 . The expectation of Ψ is $\mu := \sum_{k=1}^j (1/p_k)$, and we just need to show that $\mathbb{P}(\Psi > 1 + \kappa \sum_{k=1}^i [W_k]) \le 1/2$.

Observe that since $j \in B_i$, we have $\mu \leq \sum_{k=1}^{i} W_k \leq \sum_{k=1}^{i} [W_k]$. Hence we have

$$\mathbb{P}\Big(\Psi > 1 + \kappa \sum_{k=1}^{i} \lceil W_k \rceil \Big) \le \mathbb{P}(\Psi > \kappa \mu).$$

By Markov's inequality, this is at most $1/\kappa \leq 1/2$.

4.8 Bounding $\operatorname{Fill}_{i}^{R}(\tau_{\operatorname{end}})$

We next use the time-reversal bijection from Section 4.5 to translate Lemma 48 from a reverse randomunsticking process \tilde{R} to a random-unsticking process R by coupling the two processes.

Lemma 49. Fix λ , η and ν in (0,1) and a (λ, η, ν) -suitable send sequence **p**. Let Y be an externallyjammed process with arrival rate λ and send sequence **p**. Consider any integer $t_0 \geq j_{\min}$. Let $\tau_{end} \geq \tau_{init}$ be an integer. Then there is a coupling of

- (i) a random-unsticking process $R = R(Y, t_0)$ with
- (ii) a sample $\{Z_j^- \mid j \in bins(\tau_{end})\}$ where each Z_j^- is chosen independently from a Poisson distribution with mean $\lambda/(4p_j)$

in such a way that, for every integer $j \in B_{I(\tau_{end})-1}$, $|Fill_j^R(\tau_{end})| \ge Z_j^-$.

Proof. Fix t_0 and τ_{end} . Let Y' be an externally-jammed process with arrival rate λ and send sequence \mathbf{p} , and let $\widetilde{R} = \widetilde{R}(Y', t_0, \tau_{end})$ be a reverse random-unsticking process. Consider any $t \ge t_0 + 1$ and positive integer J. Let π be the time-reversal bijection from $\mathcal{B}^R(t_0, \tau_{end}, t, J)$ to $\mathcal{B}^{\widetilde{R}}(t_0, \tau_{end}, \tau_{end} - t + t_0 + 1, J)$, and recall from Lemma 47 that for all trajectories $B \in \mathcal{B}^R(t_0, \tau_{end}, t, J)$, we have $\mu^R(B) = \mu^{\widetilde{R}}(\pi(B))$ and $\widetilde{S}(\pi(B)) = \{\tau_{end} - t + t_0 + 1 \mid t \in S(B)\}$. From the first of these statements, the number of balls following a trajectory $B \in \mathcal{B}^R(t_0, \tau_{end}, t, J)$ in the process \widetilde{R} has the same distribution as the number of balls following the trajectory $\pi(B)$ in the process \widetilde{R} . We may therefore couple R with \widetilde{R} by identifying these trajectories, so that for every ball following trajectory B in R there is a corresponding ball following trajectory $\pi(B)$ in \widetilde{R} and vice versa.

We next show that a ball following a trajectory $B \in \mathcal{B}^R(t_0, \tau_{end}, t, J)$ contributes to $\operatorname{Fill}_j^R(\tau_{end})$ if and only if a ball following the trajectory $\pi(B)$ contributes to $\operatorname{Fill}_j^{\widetilde{R}}(\tau_{end})$, so that $|\operatorname{Fill}_j^R(\tau_{end})| = |\operatorname{Fill}_j^{\widetilde{R}}(\tau_{end})|$; given this, the lemma will follow immediately from Lemma 48. By the definition of π (Definition 46), $J(\pi(B)) = J(B)$, so (i) is the same in the definitions of $\operatorname{Fill}_j^R(\tau_{end})$ and $\operatorname{Fill}_j^{\widetilde{R}}(\tau_{end})$. Similarly, (ii) and (iii) are the same in both definitions. To see that (iv) is the same in both directions, consider $t' \in S(B)$ with $\operatorname{unstick}(B, t') = 0$. Then $\tau_{end} - t' + t_0 + 1 \in \widetilde{S}(\widetilde{B})$ by Lemma 47 and, by the definition of π , $\operatorname{unstick}(\pi(B), \tau_{end} - t' + t_0 + 1) = \operatorname{unstick}(B, t')$. Clearly, the same works in the other direction, so (iv) is the same in both definitions. Hence we do indeed have $|\operatorname{Fill}_j^R(\tau_{end})| = |\operatorname{Fill}_j^{\widetilde{R}}(\tau_{end})|$, and we are done by Lemma 48.

Finally, we show that Lemma 49 remains true when the value of $\mathcal{F}_{t_0}^R$ is exposed. (Since $bins(\tau_{end}) = B_{I(\tau_{end})-1}$ by definition, this is the only difference between the statement of Lemma 49 and that of the following lemma.)

Lemma 50. Fix λ , η and ν in (0,1) and a (λ, η, ν) -suitable send sequence **p**. Let Y be an externallyjammed process with arrival rate λ and send sequence **p**. Consider any integer $t_0 \geq j_{\min}$. Consider a random-unsticking process $R = R(Y, t_0)$. Let $F_{t_0}^R$ be any value of the filtration $\mathcal{F}_{t_0}^R$ covering the first t_0 steps of R. Let $\tau_{\text{end}} \geq \tau_{\text{init}}$ be an integer. Then there is a coupling of

- (i) The process R, conditioned on $\mathcal{F}_{t_0}^R = F_{t_0}^R$, with
- (ii) A sample $\{Z_j^- \mid j \in bins(\tau_{end})\}$ where each Z_j^- is chosen independently from a Poisson distribution with mean $\lambda/(4p_j)$

in such a way that, for every integer $j \in bins(\tau_{end})$, $|Fill_j^R(\tau_{end})| \geq Z_j^-$.

Proof. Let $i = I(\tau_{end}) - 1$. For $j \in B_i$, from its definition in Section 4.6, $\operatorname{Fill}_j^R(\tau_{end})$ is the subset of balls $\beta \in \operatorname{stuck}_j^R(t_0 + \tau_{end})$ such that $\operatorname{birth}^R(\beta) \ge \max\{t_0 + 1, t_0 + \tau_{end} - \kappa \sum_{k=1}^i [W_k]\}$. The important thing here is that $\operatorname{birth}^R(\beta) \ge t_0 + 1$.

The definition of the process R ensures that balls behave independently. Thus whether or not a ball β contributes to $\operatorname{Fill}_{j}^{R}(\tau_{\operatorname{end}})$ is independent of $\mathcal{F}_{t_{0}}^{R}$. Thus, writing R' for a copy of R conditioned on $\mathcal{F}_{t_{0}}^{R} = F_{t_{0}}^{R}$, we can couple R and R' in such a way that $\operatorname{Fill}_{j}^{R}(\tau_{\operatorname{end}}) = \operatorname{Fill}_{j}^{R'}(\tau_{\operatorname{end}})$. We conclude that Lemma 50 follows from Lemma 49.

4.9 Proving Lemma 45

Lemma 51. Fix λ , η and ν in (0, 1) and a (λ, η, ν) -suitable send sequence \mathbf{p} with $p_0 = 1$. Let Y^A and Y^B be independent externally-jammed processes with arrival rate $\lambda/2$ and send sequence \mathbf{p} and consider the two-stream externally-jammed process $T = T(Y^A, Y^B)$. Consider any integer $t_0 \geq j_{\min}$. Let $F_{t_0}^T$ be a value of $\mathcal{F}_{t_0}^T$ such that $\mathcal{E}_{init}^T(t_0)$ occurs. Let C and C' be distinct elements of $\{A, B\}$. Let $R = R(Y^C, t_0)$ be a random-unsticking process. Let $F_{t_0}^R$ be a value of $\mathcal{F}_{t_0}^R$ that is consistent with $F_{t_0}^T$. Consider any integer $\tau \geq 1$. Then there is a coupling of

- *T* conditioned on $\mathcal{F}_{t_0}^T = F_{t_0}^T$ and
- R conditioned on $\mathcal{F}_{t_0}^R = F_{t_0}^R$

such that at least one of the following occurs:

- $\mathcal{E}_{iam}^T(C', t_0, \tau)$ does not occur, or
- **unstuck**^T($C, t_0 + \tau$) \subseteq **unstuck**^R($t_0 + \tau$).

Proof. Note (from the definition of $\mathcal{F}_{t_0}^R$ and $\mathcal{F}_{t_0}^T$) since $F_{t_0}^R$ is consistent with $F_{t_0}^T$, $\mathbf{unstuck}^R(t_0) = \mathbf{unstuck}^T(C, t_0)$ and, for all $j, b_j^R(t_0) = b_j^{Y^C}(t_0)$.

We will construct the coupling step-by-step. For every positive integer τ' , we use the following notation.

- Invariant 1 for $\tau': \mathcal{E}_{iam}^T(C', t_0, \tau')$ occurs, and
- Invariant 2 for τ' : unstuck^T($C, t_0 + \tau' 1$) \subseteq unstuck^R($t_0 + \tau' 1$).

Note that the invariants for τ' only depend on steps $1, \ldots, t_0 + \tau' - 1$ of the coupled process. Our high-level strategy is as follows: Given that the coupling on these steps satisfies both invariants for τ' we will show how to extend the coupling to step $t_0 + \tau'$ to satisfy Invariant 2 for $\tau' + 1$. After that there are two possibilities.

- If $\tau' = \tau$ then we have finished, by establishing $\mathbf{unstuck}^T(C, t_0 + \tau) \subseteq \mathbf{unstuck}^R(t_0 + \tau)$.
- Otherwise, $\tau' \leq \tau 1$ and there are again two possibilities.
 - Either Invariant 1 is violated for $\tau' + 1$ in which case $\mathcal{E}_{jam}^T(C', t_0, \tau' + 1)$ does not occur so $\mathcal{E}_{jam}^T(C', t_0, \tau)$ does not occur (and we have finished),
 - or Invariant 1 is satisfied for $\tau' + 1$ and the construction of the coupling continues to the next step.

That completes the description of the high-level strategy. To finish the proof we must show that both invariants hold for $\tau' = 1$ and then we must show how to extend the coupling as promised.

We first establish the invariants for $\tau' = 1$.

- Invariant 1 follows from Equation (3) in Lemma 41.
- For Invariant 2 we wish to show $\mathbf{unstuck}^T(C, t_0) \subseteq \mathbf{unstuck}^R(t_0)$. This follows from the consistency of $F_{t_0}^T$ and $F_{t_0}^R$, as already mentioned.

To finish the proof we consider $\tau' \ge 1$ such that both invariants hold for τ' , which implies (using the definition of $\mathcal{E}_{iam}^T(C', t_0, \tau')$ in Definition 40) that

$$f(\mathbf{stuckvect}^T(C', t_0 + \tau' - 1) \ge \zeta |\operatorname{bins}(\tau' - 1)|$$

and

$$\mathbf{unstuck}^T(C, t_0 + \tau' - 1) \subseteq \mathbf{unstuck}^R(t_0 + \tau' - 1).$$

Using $\mathcal{F}_{t_0+\tau'-1}^T$ and $\mathcal{F}_{t_0+\tau'-1}^R$ we now wish to consider step $t_0 + \tau'$, and to extend the coupling to this step in such a way that $\mathbf{unstuck}^T(C, t_0 + \tau') \subseteq \mathbf{unstuck}^R(t_0 + \tau')$. Once we do this, we'll have completed the high-level strategy, hence we'll have completed the proof.

It is helpful to note by the definition of the processes that for all t and j we have

balls^T(C,t)
$$\cap$$
 $b_j^T(t) = b_j^{Y^C}(t) = b_j^R(t)$.

We will use this observation for $t = t_0 + \tau' - 1$. From this we deduce that

 $\mathbf{stucksend}^T(C,t) \subseteq \mathbf{stucksend}^R(t) \cup \mathbf{unstuck}^R(t).$

We are finally ready to do the coupling. We first determine $b'_0{}^{Y^C}(t+1) = b'_0{}^R(t+1)$ and send^{YC}(t+1) =send^R(t+1). This allows us to deduce stucksend^T(C, t+1) and stucksend^R(t+1).

We will use this to finish the coupling of step t + 1 in such a way that any ball

$$\beta \in \mathbf{stucksend}^T(C, t+1) \cap \mathbf{unstuck}^T(C, t+1)$$

is also in $\mathbf{unstuck}^R(t+1)$. Here is how we do it: If $\mathbf{stucksend}^T(C, t+1)$ is empty, there is nothing to prove. Otherwise, choose β u.a.r. from $\mathbf{stucksend}^T(C, t+1)$ (this is the random ball that may become unstuck at time t+1 in T). The probability that β is in $\mathbf{unstuck}^T(C, t+1)$ is the probability that $\mathbf{stucksend}^T(C', t+1) = \emptyset$. We will use Lemma 44 to bound this probability. First note that the lemma applies since the definition of I_0 ensures that $\zeta |B_{I_0}| \ge 4$ so $\zeta |\operatorname{bins}(\tau')| \ge 4$. By Lemma 44, the probability that $\mathbf{stucksend}^T(C', t+1) = \emptyset$ is at most

$$\exp(-\zeta|\operatorname{bins}(\tau'-1)|/16).$$

If β is in **unstuck**^{*R*}(*t*) then there is nothing to prove. Alternatively, if it is in **stucksend**^{*R*}(*t*), then the probability that it is in **unstuck**^{*R*}(*t* + 1) is $p_{\text{unstick}}(t) = \exp(-\zeta |\text{bins}(\tau' - 1)|/16)$. Thus we can extend the coupling as required, completing the proof.

It is not hard to see that the final bullet point of Lemma 51 implies $\operatorname{Fill}_{j}^{R}(\tau) \subseteq \operatorname{Fill}_{j}^{T}(C,\tau)$; this is the only difference between the statements of Lemma 51 and Corollary 52.

Corollary 52. Fix λ , η and ν in (0, 1) and a (λ, η, ν) -suitable send sequence \mathbf{p} with $p_0 = 1$. Let Y^A and Y^B be independent externally-jammed processes with arrival rate $\lambda/2$ and send sequence \mathbf{p} and consider the two-stream externally-jammed process $T = T(Y^A, Y^B)$. Consider any integer $t_0 \ge j_{\min}$. Let $F_{t_0}^T$ be a value of $\mathcal{F}_{t_0}^T$ such that $\mathcal{E}_{init}^T(t_0)$ occurs. Let C and C' be distinct elements of $\{A, B\}$. Let $R = R(Y^C, t_0)$ be a random-unsticking process. Let $F_{t_0}^R$ be a value of $\mathcal{F}_{t_0}^R$ that is consistent with $F^T(t_0)$. Consider any integer $\tau \ge 1$. There is a coupling of

- T conditioned on $\mathcal{F}_{t_0}^T = F_{t_0}^T$ and
- R conditioned on $\mathcal{F}_{t_0}^R = F_{t_0}^R$

such that at least one of the following occurs:

- $\mathcal{E}_{iam}^T(C', t_0, \tau)$ does not occur, or
- For all positive integers j, $\operatorname{Fill}_{i}^{R}(\tau) \subseteq \operatorname{Fill}_{i}^{T}(C, \tau)$.

Proof. We use the coupling of Lemma 51. If $\mathcal{E}_{jam}(C', t_0, \tau)$ does not occur, then the result is immediate. Otherwise, we have $\mathbf{unstuck}^T(C, t_0 + \tau) \subseteq \mathbf{unstuck}^R(t_0 + \tau)$. Since every ball is either stuck or unstuck, it follows that $\mathbf{stuck}^R(t_0 + \tau) \subseteq \mathbf{stuck}^T(C, t_0 + \tau)$. By the definitions of $\mathrm{Fill}_j^R(\tau)$ and $\mathrm{Fill}_j^T(C, \tau)$, it follows that $\mathrm{Fill}_i^R(\tau) \subseteq \mathrm{Fill}_j^T(C, \tau)$ as required.

Then combining the couplings of Corollary 52 (from one stream of T to R) and Lemma 50 (from R to i.i.d. Poisson variables), we get the following lemma.

Lemma 53. Fix λ , η and ν in (0,1) and a (λ, η, ν) -suitable send sequence \mathbf{p} with $p_0 = 1$. Let Y^A and Y^B be independent externally-jammed processes with arrival rate $\lambda/2$ and send sequence \mathbf{p} and consider the two-stream externally-jammed process $T = T(Y^A, Y^B)$. Consider any integer $t_0 \geq j_{\min}$. Let $F_{t_0}^T$ be a value of $\mathcal{F}_{t_0}^T$ such that $\mathcal{E}_{init}^T(t_0)$ happens. Let C and C' be distinct elements of $\{A, B\}$. Consider any integer $\tau \geq \tau_{init}$. There is a coupling of

- T conditioned on $\mathcal{F}_{t_0}^T = F_{t_0}^T$ and
- A sample $\{Z_j^- \mid j \in bins(\tau)\}$ where each Z_j^- is chosen independently from a Poisson distribution with mean $\lambda/(4p_j)$

such that at least one of the following occurs:

- $\mathcal{E}_{iam}^T(C', t_0, \tau)$ does not occur, or
- For all integers $j \in bins(\tau)$, $|Fill_i^T(C, \tau)| \ge Z_i^-$.

Finally, applying Lemma 53 twice, once with C = A and once with C = B and using the definitions of $\operatorname{Fill}_{i}^{T}(C, \tau)$ and $\mathcal{E}_{iam}^{T}(t_0, \tau)$, we get Lemma 45 as desired.

Lemma 45. Fix λ , η and ν in (0,1) and a (λ, η, ν) -suitable send sequence \mathbf{p} with $p_0 = 1$. Let Y^A and Y^B be independent externally-jammed processes with arrival rate $\lambda/2$ and send sequence \mathbf{p} and consider the two-stream externally-jammed process $T = T(Y^A, Y^B)$. Consider any integer $t_0 \geq j_{\min}$. Let $F_{t_0}^T$ be a value of $\mathcal{F}_{t_0}^T$ such that $\mathcal{E}_{init}^T(t_0)$ happens. Consider any integer $\tau \geq \tau_{init}$. Then there is a coupling of

• T conditioned on $\mathcal{F}_{t_0}^T = F_{t_0}^T$,

- a sample $\{Z_j^A \mid j \in bins(\tau)\}$ where each Z_j^A is chosen independently from a Poisson distribution with mean $\lambda/(4p_j)$, and
- a sample $\{Z_j^B \mid j \in bins(\tau)\}$ where each Z_j^B is chosen from a Poisson distribution with mean $\lambda/(4p_j)$ and these are independent of each other but not of the $\{Z_j^A\}$ values

in such a way that at least one of the following happens:

- $\mathcal{E}_{iam}^T(t_0, \tau)$ does not occur, or
- for all $j \in bins(\tau)$, $|\mathbf{stuck}_j^T(A, t_0 + \tau)| \ge Z_j^A$ and $|\mathbf{stuck}_j^T(B, t_0 + \tau)| \ge Z_j^B$.

5 The main lemma for proving Theorem 25

Our goal in this section is to prove Lemma 56, which says that if T is a two-stream externally-jammed process with suitable send sequence, and if $\mathcal{E}_{init}^T(t_0)$ occurs, then T is likely to stay t_0 -jammed forever. We start with a simple observation.

Observation 54. Suppose that **p** is (λ, η, ν) -suitable. If $i \ge I_0$ then $|\{j \in B_i \mid p_j \le p_*\}| > 2\eta |B_i|/3$.

Proof. Since **p** is (λ, η, ν) -suitable, $|\{j \in [u(i)] \mid p_j \leq p_*\}| > \eta u(i)$. At most u(i-1) of these bins are in blocks B_1, \ldots, B_{i-1} , and by definition we have $u(i-1) = u(i)/\kappa \leq \eta u(i)/3$, so at least $2\eta u(i)/3$ of these bins are in block B_i .

We will prove Lemma 56 using a union bound over all time steps. For readability, we extract the most difficult case of this bound as Lemma 55.

Lemma 55. Let **p** be (λ, η, ν) -suitable. Let *T* be a two-stream externally-jammed process with arrival rate $\lambda \in (0, 1)$ and send sequence **p**. Consider any integer $t_0 \ge j_{\min}$. Let $F_{t_0}^T$ be a value of $\mathcal{F}_{t_0}^T$ such that $\mathcal{E}_{init}^T(t_0)$ happens. Let $\tau \ge \tau_{init}$. Conditioned on $\mathcal{F}_{t_0}^T = F_{t_0}^T$, with probability at least $1 - 1/(10\tau^2)$, at least one of the following happens:

- The event $\mathcal{E}_{iam}^T(t_0, \tau)$ doesn't occur.
- The process T is t_0 -jammed for time $\tau + 1$.

Proof. Let Y^A and Y^B be independent externally-jammed processes with arrival rate $\lambda/2$ and send sequence **p** such that $T = T(Y^A, Y^B)$. Consider $C \in \{A, B\}$. Let $\{Z_j \mid j \in bins(\tau)\}$ be independent Poisson variables with $\mathbb{E}(Z_j) = \lambda/(4p_j)$. Since $\zeta = \eta\lambda/24$, by Lemma 45 and the definition of f (Definition 24),

$$\mathbb{P}\left(\overline{\mathcal{F}_{jam}^{T}(t_{0},\tau)} \text{ or } f(\mathbf{stuckvect}^{T}(C,t_{0}+\tau)) \geq \zeta |\operatorname{bins}(\tau)| \left| \mathcal{F}_{t_{0}}^{T} = F_{t_{0}}^{T} \right) \\ \geq \mathbb{P}\left(\sum_{j \in \operatorname{bins}(\tau)} p_{j}Z_{j} \geq \frac{\eta\lambda |\operatorname{bins}(\tau)|}{24}\right).$$
(13)

Since $\tau \ge \tau_{\text{init}} \ge 0$, we have $\operatorname{bins}(\tau) = B_{I(\tau)-1}$ where $I(\tau) \ge I_0 + 1$. Observation 54 ensures that there is a set $\operatorname{bins}'(\tau) \subseteq \operatorname{bins}(\tau)$ such that $|\operatorname{bins}'(\tau)| \ge 2\eta |\operatorname{bins}(\tau)|/3$ and for all $j \in \operatorname{bins}'(\tau)$, we have $p_j \le p_*$. From Equation (13) and the lower bound on $|\operatorname{bins}'(\tau)|$ we get

$$\mathbb{P}\left(\overline{\mathcal{E}_{jam}^{T}(t_{0},\tau)} \text{ or } f(\mathbf{stuckvect}^{T}(C,t_{0}+\tau)) \geq \zeta|\operatorname{bins}(\tau)| \left| \mathcal{F}_{t_{0}}^{T} = F_{t_{0}}^{T} \right) \\ \geq \mathbb{P}\left(\sum_{j\in\operatorname{bins}'(\tau)} p_{j}Z_{j} \geq \frac{\lambda|\operatorname{bins}'(\tau)|}{16}\right).$$
(14)

For each $j \in bins'(\tau)$, let 1_{Z_j} be the indicator variable of the event $Z_j \leq \lambda/(8p_j)$, so that

$$\sum_{j \in \text{bins}'(\tau)} p_j Z_j \ge \frac{\lambda}{8} \Big(|\text{bins}'(\tau)| - \sum_{j \in \text{bins}'(\tau)} \mathbb{1}_{Z_j} \Big).$$
(15)

We dominate the sum in the right-hand side of (15) above by a binomial random variable, which we will bound above using the Chernoff bound of Lemma 20. We first bound the mean μ of this variable; by the Chernoff bound of Lemma 19 applied with $\delta = 1/2$, for all $j \in bins'(\tau)$ we have

$$\mathbb{P}(1_{Z_j} = 1) = \mathbb{P}(Z_j \le \mathbb{E}(Z_j)/2) \le e^{-\mathbb{E}(Z_j)/8} \le e^{-\mathbb{E}(Z_j)/12} = e^{-\lambda/(48p_j)} \le e^{-\lambda/(48p_*)}.$$

Let $x = e^{\lambda/(48p_*)}/2$ and $\mu = e^{-\lambda/(48p_*)}|\text{bins}'(\tau)|$. Since $p_* \leq \lambda/48$, we have x > 1, which ensures the applicability of Lemma 20. From this lemma, we obtain

$$\mathbb{P}\bigg(\sum_{j\in \operatorname{bins}'(\tau)} 1_{Z_j} \ge \frac{|\operatorname{bins}'(\tau)|}{2}\bigg) \le e^{-\mu x(\log x - 1)} \le \exp\bigg(-\frac{|\operatorname{bins}'(\tau)|}{2}\bigg(\frac{\lambda}{48p_*} - 2\bigg)\bigg).$$

Recall that $|bins'(\tau)| \ge 2\eta |bins(\tau)|/3$; hence by Lemma 43, it follows that

$$\mathbb{P}\bigg(\sum_{j\in \operatorname{bins}'(\tau)} 1_{Z_j} \ge \frac{|\operatorname{bins}'(\tau)|}{2}\bigg) \le \exp\bigg(-\frac{2\eta \log(\tau)}{3 \cdot 4\kappa^2 \log(1/\nu)} \bigg(\frac{\lambda}{48p_*} - 2\bigg)\bigg).$$

The rest of the proof is simple manipulation. Since $p_* \leq \lambda/200$ we have $\lambda/(48p_*) - 2 \geq \lambda/(100p_*)$. Since $p_* \leq (\lambda \eta)/(1800\kappa^2 \log(1/\nu))$ we have $\lambda/(100p_*) \geq (18/\eta)\kappa^2 \log(1/\nu)$ so we get

$$\mathbb{P}\left(\sum_{j\in \text{bins}'(\tau)} 1_{Z_j} \ge \frac{|\text{bins}'(\tau)|}{2}\right) \le \exp(-3\log(\tau)) = \frac{1}{\tau^3}$$

Finally, since $\tau \ge \tau_{\text{init}} \ge 20$, this probability is at most $1/(20\tau^2)$. It follows from (15) that

$$\mathbb{P}\left(\sum_{j\in \text{bins}'(\tau)} p_j Z_j \ge \frac{\lambda|\text{bins}'(\tau)|}{16}\right) \ge 1 - \frac{1}{20\tau^2}$$

It follows from (14) that

$$\mathbb{P}\left(\overline{\mathcal{E}_{jam}^{T}(t_{0},\tau)} \text{ or } f(\mathbf{stuckvect}^{T}(C,t_{0}+\tau)) \geq \zeta|\operatorname{bins}(\tau)| \left| \mathcal{F}_{t_{0}}^{T} = F_{t_{0}}^{T} \right| \geq 1 - \frac{1}{20\tau^{2}}.$$

The result therefore follows by a union bound over $C \in \{A, B\}$.

Lemma 56. Let \mathbf{p} be (λ, η, ν) -suitable. Let T be a two-stream externally-jammed process with arrival rate $\lambda \in (0, 1)$ and send sequence \mathbf{p} . Consider any integer $t_0 \geq j_{\min}$. Let $F_{t_0}^T$ be a value of $\mathcal{F}_{t_0}^T$ such that $\mathcal{E}_{init}^T(t_0)$ happens. Then

$$\mathbb{P}\left(T \text{ is } t_0 \text{-jammed for all } \tau \ge 1 \mid \mathcal{F}_{t_0}^T = F_{t_0}^T\right) \ge 4/5.$$

Proof. For all $\tau \ge \tau_{\text{init}}$, let \mathcal{E}_{τ} be the event that either $\mathcal{E}_{\text{jam}}^T(t_0, \tau)$ doesn't occur or that T is t_0 -jammed for $\tau + 1$. We then have

$$\mathbb{P}\Big(\bigwedge_{\tau \ge 1} \mathcal{E}_{jam}^T(t_0, \tau) \mid \mathcal{F}_{t_0}^T = F_{t_0}^T\Big) = \mathbb{P}\Big(\mathcal{E}_{jam}^T(t_0, \tau_{init}) \land \bigwedge_{\tau \ge \tau_{init}} \mathcal{E}_{\tau} \mid \mathcal{F}_{t_0}^T = F_{t_0}^T\Big)$$

By Lemmas 41 and 55 and a union bound, we therefore have

$$\mathbb{P}\Big(\bigwedge_{\tau \ge 1} \mathcal{E}_{jam}^T(t_0, \tau) \mid \mathcal{F}_{t_0}^T = F_{t_0}^T\Big) \ge \frac{49}{50} - \sum_{\tau \ge \tau_{init}} \frac{1}{10\tau^2} \ge \frac{49}{50} - \frac{\pi^2}{60} \ge \frac{4}{5},$$

as required.

_		

6 **Proof of Main Theorems**

Theorem 25. Fix λ , η and ν in (0, 1). Let **p** be a (λ, η, ν) -suitable send sequence. Let T be a two-stream externally-jammed process with arrival rate λ and send sequence **p**. Let $\mathcal{R}^T = \{t \mid \mathbf{stuck}^T(t) = \emptyset\}$. With probability 1, \mathcal{R}^T is finite.

Proof. For every positive integer i, let R_i be the i'th smallest time step in \mathcal{R}^T . If $|\mathcal{R}^T| < i$, then we define $R_i = \infty$. Since $R_1 = 0$, for all $i \ge 2$ we have

$$\mathbb{P}(R_i < \infty) = \prod_{k=2}^{i} \mathbb{P}(R_i < \infty \mid R_{i-1} < \infty).$$
(16)

Let \mathcal{E}_i be the event that $R_{i-1} < \infty$ and that $\mathcal{E}_{init}^T(R_{i-1} + t)$ occurs for some positive integer t such that $R_{i-1} + t < R_i$. Observe that conditioned on $R_{i-1} < \infty$, the probability of \mathcal{E}_i occurring does not depend on i since stuck^T evolves independently of unstuck balls; let $p_{init} = \mathbb{P}(\mathcal{E}_i \mid R_{i-1} < \infty)$. Then by (16), we have

$$\mathbb{P}(R_i < \infty) \le \prod_{k=2}^{i} \left(1 - p_{\text{init}} + p_{\text{init}} \mathbb{P}(R_i < \infty \mid \mathcal{E}_i, R_{i-1} < \infty) \right).$$
(17)

Observe that if $\mathcal{E}_{\text{init}}(R_{i-1}+t)$ occurs for some $t < R_i - R_{i-1}$, and T is $(R_{i-1}+t)$ -jammed for all $\tau \ge 1$, then we have we must have stuck^T $(R_{i-1}+u) \ne \emptyset$ for all $u \ge 1$; hence $R_i = \infty$.

It follows by applying Lemma 56 to (17) that

$$\mathbb{P}(R_i < \infty) \le \prod_{k=2}^{i} \left(1 - p_{\text{init}} + p_{\text{init}}/5 \right) \le e^{-4(i-1)p_{\text{init}}/5}.$$
(18)

By (18), we have $\sum_i \mathbb{P}(R_i < \infty) < \infty$. It follows by the Borel-Cantelli lemma that with probability 1 we have $R_i = \infty$ for all but finitely many values of *i*, which proves the theorem.

Corollary 26. Fix λ , η and ν in (0, 1). Let **p** be a (λ, η, ν) -suitable send sequence. Let X be a backoff process with arrival rate λ and send sequence **p**. Let $\mathcal{R}^X = \{t \mid \mathbf{balls}^X(t) = \emptyset\}$. Then, with probability 1, \mathcal{R}^X is finite. Hence, X is transient so it is unstable.

Proof. Let T be a two-stream externally-jammed process with arrival rate λ and send sequence \mathbf{p} . Let $\mathcal{R}^T = \{t \mid \mathbf{stuck}^T(t) = \emptyset\}$. Theorem 25 shows that, with probability 1, \mathcal{R}^T is finite. Let Y be an externally-jammed process with arrival rate λ and send sequence \mathbf{p} , and let $\mathcal{R}^Y = \{t \mid \mathbf{stuck}^Y(t) = \emptyset\}$. By the coupling of Y and T from Lemma 23, with probability 1, \mathcal{R}^Y is finite. Finally, by the coupling of X and Y from Observation 22, it follows that with probability 1, \mathcal{R}^X is finite.

In order to make the proof of Theorem 5 more convenient, we first re-state the definition of LCED.

Definition 4. A send sequence \mathbf{p} is LCED ("largely constant with exponential decay") if it satisfies the following properties:

- (*i*) "Largely constant": For all $\eta > 0$, there exists c > 0 such that for infinitely many n, $|\{j \le n : p_j > c\}| \ge (1 \eta)n$.
- (*ii*) "with exponential decay": **p** has an infinite subsequence $(p_{\ell_1}, p_{\ell_2}, ...)$ which satisfies $\log(1/p_{\ell_x}) = \Theta(\ell_x)$ as $x \to \infty$.
- (*iii*) "(but without super-exponential decay)": $\log(1/p_j) = O(j)$ as $j \to \infty$.

Theorem 5. Let \mathbf{p} be a send sequence which is not LCED. Then for every $\lambda \in (0,1)$ the backoff process with arrival rate λ and send sequence \mathbf{p} is unstable.

Proof. Suppose that **p** is not LCED. Fix $\lambda \in (0, 1)$. Let X be a backoff process with arrival rate λ and send sequence **p**. We split the analysis into cases depending on which case of Definition 4 fails.

Case 1 (iii) fails: By definition, since (iii) fails, there is no real number C > 0 such that, for all but finitely many j, $\log(1/p_j) \le Cj$. So for all C > 0, there are infinitely many j such that $\log(1/p_j) > Cj$, which implies $p_j < \exp(-Cj)$. Taking $C = \log(2/(\lambda p_0))$, we conclude that there are infinitely many j such that $p_j < (\lambda p_0/2)^j$. Lemma 7 implies that X is unstable.

Case 2 (iii) holds but (i) fails: By definition, since (i) fails, there is an $\eta > 0$ such that for all c > 0, for all but finitely many n, $|\{j \le n : p_j \le c\}| > \eta n$. Taking $c = p_*(\lambda p_0, \eta, \nu)$ for positive ν , we conclude that there is an $\eta > 0$ such that for all $\nu > 0$ and for all but finitely many n, $|\{j \le n : p_j \le p_*(\lambda p_0, \eta, \nu)\}| > \eta n$.

Let \mathbf{p}' be the send sequence derived from \mathbf{p} by setting $p'_0 = 1$ and, for every positive integer j, $p'_j = p_j$. Given the $\eta > 0$ that we have just identified, our goal will be find $\nu > 0$ such that \mathbf{p}' is $(\lambda p_0, \eta, \nu)$ -suitable (Definition 14). To achieve this, we need to show that we can choose $\nu > 0$ such that, for all but finitely many n, $\nu^n < p_n = p'_n$.

Here we will use the fact that (iii) holds. Since $\log(1/p_j) = O(j)$, there exists C > 0 such that for all but finitely many j, $1/p_j < e^{Cj}$. So taking $\nu = e^{-C} < 1$ we have that for all but finitely many j, $p_j > \nu^j$. So we have now shown that \mathbf{p}' is $(\lambda p_0, \eta, \nu)$ -suitable.

Let X' be a backoff process with arrival rate λp_0 and send sequence p'. Let $\mathcal{R}^{X'} = \{t \mid \mathbf{balls}^{X'}(t) = \emptyset\}$. By Corollary 26, with probability 1, $\mathcal{R}^{X'}$ is finite.

Finally, let X be a backoff process with arrival rate λ and send sequence **p**. Let $\mathcal{R}^X = \{t \mid \mathbf{balls}^X(t) = \emptyset\}$. By the coupling of X and X' from Observation 21, with probability 1, \mathcal{R}^X is finite. Hence, the process X is transient, so it is unstable.

Case 3 (iii) holds but (ii) fails: In this case $\log(1/p_j) = o(j)$ so the result follows from Corollary 9. Here is an argument that $\log(1/p_j) = o(j)$. Since (iii) holds but (ii) fails, there is no infinite subsequence $(p_{\ell_1}, p_{\ell_2}, ...)$ that satisfies $\log(1/p_{\ell_x}) = \Omega(\ell_x)$. Thus for every infinite subsequence $(p_{\ell_1}, p_{\ell_2}, ...)$ and for every C > 0, infinitely many n satisfy $\log(1/p_{\ell_n}) < C\ell_n$. This implies that for all C > 0 and all infinite subsequences $(p_{\ell_1}, p_{\ell_2}, ...)$, some n has $\log(1/p_{\ell_n}) < C\ell_n$. Hence, for all C > 0 there is no infinite subsequence $(p_{\ell_1}, p_{\ell_2}, ...)$ that satisfies $\log(1/p_{\ell_n}) < C\ell_n$. Hence, for all C > 0 there is no infinite subsequence $(p_{\ell_1}, p_{\ell_2}, ...)$ that satisfies $\log(1/p_{\ell_n}) \ge C\ell_x$. Finally, for all C > 0, there is a j_C such that for all $j \ge j_C$, $\log(1/p_j) < Cj$ so $\log(1/p_j) = o(j)$, as required.

We can now use Theorem 5 to derive the other theorems stated in the introduction.

Theorem 2. For every $\lambda \in (0, 1)$ and every monotonically non-increasing send sequence $\mathbf{p} = p_0, p_1, \ldots$, the backoff process with arrival rate λ and send sequence \mathbf{p} is unstable.

Proof. We prove that \mathbf{p} is not LCED and use Theorem 5. Suppose for contradiction that \mathbf{p} is LCED. Let c > 0 be such that $p_j \ge c$ for infinitely many values of j; by (i) of Definition 4 applied with any $\eta > 0$, such a c must exist. Let $(p_{\ell_1}, p_{\ell_2}, \dots)$ be the infinite exponentially-decaying subsequence of \mathbf{p} guaranteed by (ii), and observe that there exists i_c such that $p_{\ell_i} < c$ for all $i \ge i_c$. Since \mathbf{p} is monotonically non-increasing, it follows that for all $j \ge i_c$ we have $p_j \le p_{i_c} \le c$, giving a contradiction.

Theorem 3. Let \mathbf{p} be a send sequence. Let $m_{\mathbf{p}}(n)$ be the median of p_0, \ldots, p_n . Suppose that $m_{\mathbf{p}}(n) = o(1)$. Then for every $\lambda \in (0, 1)$ the backoff process with arrival rate λ and send sequence \mathbf{p} is unstable.

Proof. We observe that **p** is not LCED and use Theorem 5. In particular, we will observe that item (i) in the definition of LCED cannot be satisfied for $\eta = 1/2$. Item (i) implies that there is a c > 0 such that for infinitely many $n, m_{\mathbf{p}}(n) > c$. But this is inconsistent with $m_{\mathbf{p}}(n) = o(1)$.

References

- IEEE standard for information technology-telecommunications and information exchange between systems local and metropolitan area networks — specific requirements — part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications. *IEEE Standard 802.11-2016* (*Revision of IEEE Std 802.11-2012*, pages 1–3534, 2016.
- [2] Norman Abramson. THE ALOHA SYSTEM. In Norman Abramson and Franklin Kuo, editors, *Computer-Communication Networks*, pages 281–285. Prentice-Hall, 1973.
- [3] Hesham Al-Ammal, Leslie Ann Goldberg, and Philip D. MacKenzie. An improved stability bound for binary exponential backoff. *Theory of Computing Systems*, 34(3):229–244, 2001.
- [4] David Aldous. Ultimate instability of exponential back-off protocol for acknowledgment-based transmission control of random access communication channels. *IEEE Transactions on Information Theory*, 33(2):219–223, 1987.
- [5] Mark Allman, Vern Paxson, and Ethan Blanton. TCP congestion control. RFC, 5681:1–18, 2009.
- [6] Michael A. Bender, Tsvi Kopelowitz, William Kuszmaul, and Seth Pettie. Contention resolution without collision detection. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC*, pages 105–118. ACM, 2020.
- [7] Haimin Chen, Yonggang Jiang, and Chaodong Zheng. Tight trade-off in contention resolution without collision detection. In Avery Miller, Keren Censor-Hillel, and Janne H. Korhonen, editors, PODC '21: ACM Symposium on Principles of Distributed Computing, pages 139–149. ACM, 2021.
- [8] Bogdan S. Chlebus. Randomized communication in radio networks. *CoRR*, abs/1801.00074, 2018. URL: http://arxiv.org/abs/1801.00074.
- [9] AWS documentation. Error retries and exponential backoff in AWS. URL: https://docs.aws.amazon.com/general/latest/gr/api-retries.html.
- [10] Google Drive documentation. Handle errors and improve performance. URL: https://developers.google.com/drive/api/guides/handle-errors?hl=en#exponential-
- [11] Anthony Ephremides and Bruce Hajek. Information theory and communication networks: an unconsummated union. *IEEE Transactions on Information Theory*, 44(6):2416–2434, 1998.
- [12] Leslie Ann Goldberg, Mark Jerrum, Sampath Kannan, and Mike Paterson. A bound on the capacity of backoff and acknowledgment-based protocols. *SIAM Journal on Computing*, 33(2):313–331, 2004.
- [13] Leslie Ann Goldberg, Philip D. MacKenzie, Mike Paterson, and Aravind Srinivasan. Contention resolution with constant expected delay. *Journal of the ACM*, 47(6):1048–1096, 2000.
- [14] Jonathan Goodman, Albert G. Greenberg, Neal Madras, and Peter March. Stability of binary exponential backoff. *Journal of the ACM*, 35(3):579–602, 1988.
- [15] Albert G. Greenberg, Philippe Flajolet, and Richard E. Ladner. Estimating the multiplicities of conflicts to speed their resolution in multiple access channels. *Journal of the ACM*, 34(2):289–325, 1987.
- [16] Johan Håstad, Frank T. Leighton, and Brian Rogoff. Analysis of backoff protocols for multiple access channels. SIAM Journal on Computing, 25(4):740–774, 1996.

- [17] Maurice Herlihy and J. Eliot B. Moss. Transactional memory: Architectural support for lock-free data structures. In *Proceedings of the 20th Annual International Symposium on Computer Architecture, San Diego, CA, USA, May 1993*, pages 289–300. ACM, 1993.
- [18] Frank P. Kelly and Iain M. MacPhee. The number of packets transmitted by collision detect random access schemes. *Annals of Probability*, 15:1557–1568, 1987.
- [19] Robert Metcalfe and David R. Boggs. Ethernet: Distributed packet switching for local computer networks. *Communications of the ACM*, 19(7):395–404, 1976.
- [20] Michael Mitzenmacher and Eli Upfal. Probability and Computing: Randomized Algorithms and Probabilistic Analysis. Cambridge University Press, 2005.
- [21] Jeannine Mosely and Pierre Humblet. A class of efficient contention resolution algorithms for multiple access channels. *IEEE Transactions on Communications*, 33(2):145–151, 1985.
- [22] Prabhakar Raghavan and Eli Upfal. Stochastic contention resolution with short delays. *SIAM Journal on Computing*, 28(2):709–719, 1998.
- [23] Ravi Rajwar and James R. Goodman. Speculative lock elision: enabling highly concurrent multithreaded execution. In Yale N. Patt, Josh Fisher, Paolo Faraboschi, and Kevin Skadron, editors, Proceedings of the 34th Annual International Symposium on Microarchitecture, Austin, Texas, USA, December 1-5, 2001, pages 294–305. ACM, 2001.
- [24] Lawrence G. Roberts. ALOHA packet system with and without slots and capture. ACM SIGCOMM Computer Communication Review, 5(2):28–42, 1975.
- [25] Devavrat Shah and Jinwoo Shin. Randomized scheduling algorithm for queueing networks. *Annals of Applied Probability*, 22:128–171, 2012.
- [26] Devavrat Shah, Jinwoo Shin, and Prasad Tetali. Medium access using queues. In IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, pages 698–707. IEEE Computer Society, 2011.